FINAL REPORT
ARTIFICIAL INTELLIGENCE (AI) CENTER OF EXCELLENCE
AT
THE UNIVERSITY OF PENNSYLVANIA

Army Research Office, Grant No.: DAAL03-89-C-0031

July 1995

Approved for public release; distribution unlimited.

THE VIEW, OPINIONS, AND/OR FINDINGS CONTAINED IN THIS REPORT ARE THOSE OF THE AUTHOR(S) AND SHOULD NOT BE CONSTRUED AS AN OFFICIAL DEPARTMENT OF THE ARMY POSITION, POLICY, OR DECISION, UNLESS SO DESIGNATED BY OTHER DOCUMENTATION.

Contents

1 Statement of the Problem Studied

The ARO AI Research and Education Grant was initially awarded to the University of Pennsylvania in June 1984, and renewed in October, 1989. The AI Center was established to integrate and coordinate our various research efforts in AI in which we have excellent strengths. The key thrusts were in the following areas: 1) Natural language processing: language and speech; 2) Machine perception and robotics: exploration and perceptual development; 3) Task oriented computer animation; 4) Programming structures for databases and knowledge bases; 5) Parallel processing in Artificial Intelligence.

2 Summary of the Most Important Results

- **Computer Graphics and Animation**

  A long range goal has been to extend the capabilities of their human figure modelling software, *Jack*, into highly interactive, yet reasonably "behaved" models. The major advances are as follows: We have completed a series of real-time, interactive strength display programs running in Jack. These graph end-effector force or joint torque in any body posture, track maximum strength for a limb, show comfort level, work, or energy expenditure, and spatially display acceptable and unacceptable force level regions for a limb holding a weight. Another advance is the extension of our strength-guided motion work into coordinated motion for two hands or two people. The third area is the extension of the inverse kinematics constraint solver to include the center of mass of a figure as the subject of a constraint. This permits "balanced reach": the constraint of center of mass to lie above a support polygon, and automatic posture adjustment to maintain such a constraint during interactive figure manipulation. The fourth area is the beginning of general clothing development for human figure models, and especially the clothing of actual contour-scanned bodies. Finally, we have developed a flexible 17-link human torso model that has amazingly lifelike motion capabilities and limitations. We integrated these better body models into the distribution versions of the *Jack* software.

  We have developed and expanded our research software, *Jack*, for articulated figure animation and human factors evaluation. New features include new Army ANSUR anthropometry data, better spreadsheet functionality in selecting human body sizes, full integration of the biostereometric high resolution realistic body models, compliant hand grasp behaviors, distributed system networking for multiple person or object control, and interfaces with Ascension Flock of Birds 6-axis sensing devices for direct human operator ("Virtual Reality") input. The walking algorithm was improved with actual biomechanical data. Approach and repel behaviors were implemented in order to study higher level behavioral simulation. Parallel algorithms for motion planning (as a computational geometry problem) are being investigated.

  Jack '94 (v.5.8) was released in March, 1994. Additions and enhancements to *Jack* include support for an improved human body geometry based on detailed Viewpoint DataLabs human models, a posture control network for transitioning the figure form one pose to another (e.g., such a s standing to prone) possibly passing through intermediate poses (e.g. kneel) on the way, improvements in collision avoidance to accommodate strength data, and real-time locomotion which responds to external force loads by changing body posture or gait. We have also built novel graphics rendering software for the Silicon Graphics systems that does not require special hardware; it does real-time shadows, reflections, and refractions through existing pipelined instructions.

- **Natural Language Processing**

  The Combinatory Categorial Grammar theory of intonation, discourse information and syntax has been further extended, in particular in respect of the theory of discourse information where the system is being extended to cover further distinctions of discourse meaning. So far, these categories include topic (or what is being talked about), comment (or what is being added to previously established information), and focus (or what information stands in contrast). The first phase of the practical application of CCG to the synthesis of contextually appropriate intonation in spoken language responses to database queries is complete. We are now extending the system in application to the Penn-based TraumAID medical expert system. Work is continuing on the implementation of a theory of tense and aspect using an event calculus, which is also required in this domain.

  We have built a large scale lexicalized tree-adjoining grammar (LTAG) parser and lexicon. The lexicon has over 200,000 items. We have also built a morphological analyzer and a morphological database. The parser is being tested on large scale corpora, *IBM* Manuals, and *The Wall Street Journal.*

  We have developed new techniques for applying part-of-speech disambiguities techniques to large tree fragments, where these fragments are the trees of the lexicalized tree-adjoining grammars (LTAG). These techniques lead to about 77% success rate on a test corpus of the *Wall Street Journal* (WSJ). Success here means that this technique assigns the same tree fragment to a word in a sentence that the LTAG would have assigned, if the sentence was parsed without the disambiguation phase.

  We have extended our word-sense disambiguation technique to the automatic restoration of word accents in French and Spanish with 99% overall accuracy.

  A new spoken reply systems (IBIS) assigns intonation contours based on contrasts derived from a simple data-base model, correcting many cases of mis-synthesis by the current dominant algorithm.

- **Diagnostic Reasoning and Expert Systems**

  We have focused on the development of efficient algorithms for diagnostic reasoning and therapy planning over time. In one research project, a graphical decision-theoretic framework, called an Influence Diagram (ID), is used. We have been able to extend the ID to incorporate time explicitly. This temporal ID has been applied to the diagnosis and treatment of acute abdominal pain, a domain in which the temporal aspects are important to accurate diagnosis and successful therapy. In particular, we have developed criteria for conducting efficient reasoning for this domain (the general task is NP-complete) through graph-reduction techniques. Statistical criteria for loss of decision-making accuracy can be specified for this graph-reduction process.

  We have also developed a mixed qualitative/quantitative simulation system, called QobiSIM. QobiSIM has been used to develop a cardiovascular simulation to be incorporated into the TraumAID system. This cardiovascular model can accept a mixture of exact values (such as blood pressure and heart rate) along with estimates of blood loss, and is being designed to produce predictions for required infusion rates, time to shock (or death), etc. Such predictions will assist the TraumAID system in therapy planning.

  We have focused on the analysis of the results of the large-scale retrospective evaluation of TraumAID 2.0, a medical decision system developed by us. Results that emerged from the three judges (expert trauma surgeons) carrying out a blind evaluation of the three versions

each of 100 cases of penetrating abdomen and/or chest trauma was that TraumAID was found preferable to actual care by a ratio of 64:17 and to TraumAID 1.0 by a ratio of 62.9. These results are statistically significant.

We have continued our work on the graphical interface for use by Emergency Room nurse-scribes that will enable us to perform prospective evaluation of TraumAID and get to grips with the human factors problem of integrating it into clinical practice. The delivery mode we have developed is real-time critiquing of physician orders. In this mode, the system reserves comment unless there is a clinically significant difference between the physician's orders and what TraumAID has concluded to be the most important management goals and procedures. Further work will involve (1) quantifying significant differences, (2) studying interactions among trauma team members, focusing on cases where one member's proposed actions are rejected or criticized by another, and (3) installing TraumAID in the MCP Emergency Room.

- **Computer Vision and Robotics**

Many tasks in active perception require the ability to combine different information from a variety of sensors that relate to one or more features of the environment. The purpose of this research is to examine sensor fusion problems for linear location data models using statistical decision theory (SDT). The contribution of this research is the application of SDT to obtain: (i) a robust test of the hypothesis that data from different sensors are consistent; and (ii) a robust procedure for combining the data that pass this preliminary consistency test. Here, robustness refers to the statistical effectiveness of the decision rules when the probability distributions of the observation noise and the a priori position information associated with the individual sensors are uncertain. The standard linear location data model refers to observations of the form: $Z = \theta + V$, where V represent additive sensor noise and θ denotes the "sensed" parameter of interest to the observer. While the theory addressed in this research applies to many uncertainty classes, the primary focus of this research is on asymmetric and/or multimodal models, that allow one to account for very general deviations from nominal sampling distribution.

We have extended our theoretical model of Cooperative Behaviors which uses the formalism of Discrete Event Systems (DES) into shared events. This is necessary when more than one sensor is controlling a vehicle and there may be conflicting directives (i.e., when one says, 'go to the right' and the other, 'go to the left'). This model has been tested in an experimental setting when a vehicle is tracking a target while also avoiding an obstacle. Secondly, we have completed a system which has given a digitized free form part we can flexibly fit spline functions depending upon the local curvature.

In the work of understanding shadows we have used color for making an hypothesis about which patches on a scene may be shadows. Then, from range data, some a priori information about the illumination source, or by probing the source of light, we can determine which of the hypothetical shadow patches are true shadows due to obscuration and which are not shadows but are due to other illumination effects.

We have extended our previously developed technique for the identification of unoccluded moving human limb outlines to fully segment moving human outlines. Through this technique we can also recover the shape and motion of each identified part. Secondly, we developed new techniques for the integration of multiple viewpoints and tracking of multiple independently moving objects using a combination of aspect-based and physics-based modeling techniques. We have also extended our algorithm for the shape and motion estimation of the heart's left

ventricle by making it more computationally efficient and by applying it to various patient data in order to quantify normal and abnormal behavior.

We have developed a method for the efficient re-rendering of a graphical scene under a directional illuminant at an arbitrary orientation. This re-rendering is accomplished via linear combination of a set of pre-rendered "basis" images. The technique can be used, for example, to animate an architectural scene under changing illumination conditions (e.g., sun moving across sky, sun going behind clouds).

Many multi-dimensional signal processing problems require the computation of signal gradients or directional derivatives. Traditional derivative estimates based on adjacent or central differences are often inappropriate for multi-dimensional problems. As replacements for these traditional operators, we have designed a set of matched pairs of derivative filters and low-pass prefilters. We have demonstrated the superiority of these filters over simple difference operators.

- **Neurally Motivated Models**

We have developed a neurally motivated model of reflexive (rapid) reasoning. We prepared a detailed response to over thirty commentaries written by experts in the fields of AI, psychology, philosophy, and biology on a major article describing our system. The article, commentaries, and our response has appeared in the interdisciplinary journal, *Behavioral and Brain Sciences*, (1993). We have designed a related system that would learn new rules with the appropriate semantics restrictions. We have also carried out a preliminary design for implementing the reasoning system on a CM-2.

We have developed a connectionist system for recognizing handwritten words. The system consists of three interacting modules, namely, i) a coarse character recognition/segmentation character boundaries based on a coarse form of recognition, ii) a fine grained isolated character recognition system that scans the image along several axes and recognizes the resulting spatiotemporal pattern, and iii) a procedural controller that integrates and controls the activity of the two recognizers. We have completed the design of the coarse and fine grained digit recognizer. In particular, we have i) retrained the digit recognizers to deal with overlapping and partial digits and ii) retrained and tested these modules using a large data set made available by NIT. Work on the design of the procedural controller is in progress. Currently, the system achieves 99% accuracy with a rejection rate of only 9.5%.

- **Programming Language Design and Theory**

We have been using Milner's recent calculus, the pi-calculus, to specify the operational semantics of programming languages, particular logic and functional programming languages. Viewing various aspects of programming language implementations using processes is an exciting and fruitful new abstraction. For example, it is natural to consider a reference cell within a functional programming language as a mini-process separate from the main evaluation thread: reading and setting a reference can be done by synchronizing with the reference cell. Since the semantics of references has proved to be difficult, this approach towards specification allow us to employ notions of semantics from concurrency to this problem. The process paradigm again helps to provide an exciting and interesting specification of the complex control structure of Prolog: the mobility feature of the pi-calculus made it possible to model the use of pointers and circular data structures that are found in Prolog. The resulting specification is modular and clear, something less true about other attempts at such specification.

We have developed a theory of a new specification logic called Forum. Forum is a presentation of all of linear logic that can be viewed as a logic programming language that modularly extends Prolog and lambda Prolog. Forum extends these other languages in that it allows for the specification of concurrent computations via its generalization of the multiset rewriting paradigm. To understand this use of linear logic as a specification language, we have used Forum to specify various problems in the semantics of computational systems. For example, we have managed to give declarative and modular specifications of such high-level imperative programming features as side-effects, exceptions, and continuation-passing. While simple in this setting, providing formal and modular specifications has proved difficult in other frameworks, such as those based on denotational semantics and structured operational semantics.

We have continued our work in the area of programming languages and in the area of artificial intelligence. The first involved developing an approach to describing the semantics of functional programs at a level of abstraction capable of expressing sharing, copying, and value liveness. Programmers need this information to reason about the resource usage of programs, but knowing how the compiler for the language is written demands too much knowledge of detail; hence having an abstract model that describes only what is wanted will be useful. We have also worked on the development of a collection of primitive operations on *anti-chains,* an order-theoretic concept which can be viewed as basic to several algorithm in knowledge representation and machine learning. The discovery of the right primitives will make it possible to share ideas (and code) between these diverse applications.

- **Heterogeneous Databases**

The problem area that we have been addressing is providing integrated access to multiple, pre-existing databases that are connected over a network. We have continued to develop a general technique for merging database schemas that has a number of advantages over existing techniques, the most important of which is that schemas are placed in a partial order that has bounded joins. This means that the merging operation, when it succeeds, is both associative and commutative, i.e. that the merge of schemas is independent of the order in which they are considered– a property not possessed by existing methods. The technique is interactive in that users make assertions about the relationships between the nodes of the schemas to be merged. These assertions are then considered to be elementary schemas, and are combined with the schemas using precisely the same merging operation. The technique is general and can be applied to a variety of data models. It can also deal with certain cardinality constraints that arise through the imposition of keys. A prototype implementation, together with a graphical interface, has been developed.

Formal techniques have been developed for heterogeneous databases, and that there is considerable practical benefit to be gained from this. To that end, work has been initiated in in related areas:

1.) Schema and Data Integration: Developed techniques for providing a uniform view of multiple databases, which are possibly cast in different data models. Also developed techniques for extracting data from the multiple sources, which may be incomplete or contradictory.

2.) Database Query languages: Exploring languages that will address many of the issues faced in data integration. For example, the language should uniformly handle a variety of data models as well and have greater expressive power that those in, say, the tradition of SQL and its descendants. This allows integrators to capture data from several sources, express the integration relationships and create an integrated instance.

These techniques will be used in an effort to develop tools for "transforming" schemas. In many applications that are being developed, the databases are constantly "evolving", not only in information content (i.e. at the data level) but in structure (i.e. at the schema level). Tools are being developed to allow the designers to assure consistency of design, and graceful upgrading of applications.

- **Real Time Systems**

  A fundamental issue in the automated analysis of concurrent systems is the efficient generation of the reachable state space. Since it is not possible to explore all the reachable states of a system if the number of states is very large or infinite, we have developed a technique for minimizing the state space. This approach is to cluster subsets of states into equivalent classes. We assume that concurrent systems are specified as communicating state machines with arbitrary data space. We describe a procedure for constructing a minimal reachability state graph from communicating state machines. As an illustration of our approach, we analyze a producer-consumer program written in Ada.

  There is a growing interest in describing the timing properties and resource constraints of pipelined superscalar processor instructions at a high level. Superscalar processors can issue and execute multiple instructions simultaneously. The degree of parallelism depends on the multiplicity of hardware functional units as well as data dependencies among instructions. Thus, the timing properties of a superscalar program are difficult to analyze and predict. On the other hand, we have been developing a real-time process algebra, called ACSR, for specification and analysis of timed concurrent systems. We have developed techniques for modelling the instruction-level architecture of a superscalar processor using ACSR and for deriving the temporal behavior of an assembly program using the ACSR laws. There are several advantages to our approach. First, the translation of instructions to ACSR processes provides formal semantics and also precise documents on the timing and resource constraints of superscalar instructions. Second, since the ACSR interpreter already exists, the development of a timing analyzer for a new superscalar processor requires only a simple translator from instructions to processes. Third, it is possible to verify and synthesize hardware and software of embedded systems since both can be specified in the same framework.

  As computers become ubiquitous, they are increasingly used in safety critical environments. Since many safety critical applications are real-time systems, automated analysis technique of real-time properties is desirable. Most widely used automated analysis techniques are based on state space exploration. Automatic analysis techniques based on state space exploration suffer from the state space explosion problem. In particular, a real-time system may have an unbounded number of states due to infinitely many possible time values. We have developed an approach for generating a finite and efficient representation of the reachable states called a timed reachability graph for a real-time system. In this work, a real-time system is specified using a timed automaton which is a timed extension of the well-known finite automaton. Our approach for coping with the state explosion problem is to extract timing information from states and to represent it as relative time relations between transitions. We have also developed an algorithm for computing the minimum and maximum time bounds between executions of two actions from a timed reachability graph to determine timing properties.

- **Algorithms in Computational Biology**

  We have focused on algorithmic problems concerned with reconstructing evolutionary trees. We have made progress in two problem areas during this period: (1) the inference of evolu-

tionary history for species sets from character data, (2) reconcilation of a set of evolutionary trees.

For the problem of inferring evolutionary history, we have found an $O(2^{2r}nk^2)$ algorithm for the Perfect Phylogeny Problem, a classical problem in this area, concerned with the inference of evolutionary trees from k r-state characters defining n species. This algorithm is particularly useful for species sets defined by biomolecular data, or other data sets in which the number of character states can be bounded. We have found an $O(kn^{0.5})$ algorithm for the reconstruction of rooted trees from subtrees, improving on the previous best $O(kn)$ algorithm.

3 List of Publications: October 1989 - March 1995

Please see attached Appendix A.

4 List of all Participating Scientific Personnel showing any advanced degrees earned by them while employed on the project

Ph.D. Graduates (AI and AI Related) (1989 — present)

GRAD	NAME	DISSERTATION TOPIC	FIRST POSITION
Ph.D. Recipients 1989			
Dec 89	**Amy Felty**	"Specifying and Implementing Theorem Provers in a Higher-Order Logic Programming Language" Advisor: *Miller*	Visiting researcher, INRIA, Sophia Antipolis, France
GRAD	**NAME**	**DISSERTATION TOPIC**	**FIRST POSITION**
Ph.D. Recipients 1990			
May 90	**Sunil Shende**	"Computational Feasibility of Some Constrained Grammatical Formalisms for Natural Language" Advisor: *Palis*	Assistant Professor, University of Nebraska
Aug 90	**Yves Schabes**	"Mathematical and Computational Aspects of Lexicalized Grammars" Advisor: *Joshi*	Research Associate, Department of Computer Science, University of Pennsylvania
Aug 90	**Isaac Rudomin**	"Simulating Cloth Using a Mixed Geometric-Physical Method" Advisor: *Badler*	Monterrey Institute of Technology, Mexico
Dec 90	**Jugal Kalita**	"Natural Language Control of Animation of Task Performance in A Physical Domain" Advisor: *Badler*	Assistant Professor, University of Colorado, Colorado Springs, CO

GRAD	NAME	DISSERTATION TOPIC	FIRST POSITION
Ph.D. Recipients 1991			
May 91	**Tarek Almedin**	"Robotic manipulation Using A Behavioral Framework" Advisor *Badler*	Asst. Professor, Computer Science, California State Univ. Fresno, CA
May 91	**Janez Funda**	"Teleprogramming: Toward Delay-Invariant Remote Manipulation" Advisor: *Paul*	IBM, Thomas J. Watson Research Center, Yorktown Heights, NY
May 91	**Alok Gupta**	"Volumetric Segmentation of Complex 3D Scenes Using Parametric Shape Models" Advisor: *Bajcsy*	Post Doctoral Research, University of Pennsylvania, General Robotics and Active Sensory Perception Laboratory, Philadelphia, PA
May 91	**David Wei**	"Fast Parallel Routing and Computation On Interconnection Networks" Advisor: *Badler*	Asst. Professor, Computer Science, Radford University, Radford, VA
Aug 91	**Sanjay Agrawal**	"Robotic Manipulation Using A Behavioral Framework" Advisor: *Bajcsy*	Microsoft Corporation, Redmond, WA
Aug 91	**Richard Gerber**	"Communicating Shared Resources: A Model for Distributed Real-Time Systems" Advisor: *Lee*	Asst. Professor, Computer Science, University of Maryland, College Park, MD
Aug 91	**Victor Wolfe**	"Supporting Real-Time Concurrency" Advisor: *Davidson/Lee*	Asst. Professor, Computer Science, University of Rhode Island, Kingston, RI
Dec 91	**Brant Cheikes**	"A New Perspective in Cooperative Question Answering" Advisor: *Webber*	Post-doctoral research, University of Linkoping, Department of Computer and Information Science, Linkoping, Sweden
Dec 91	**Catherine Pelachaud**	"Communication and Coarticulation" Advisor: *Badler*	Research Staff, Computer Science, University of Roma, La Sapienza, Italy
Dec 91	**Cary Bradford Phillips**	"Interactive Postural Control of Articulated Geometric Figures" Advisor: *Badler*	Pacific Data Images, Sunnyvale, CA
Dec 91	**Dennis G. Shea**	"Parallel Processing: Experimentation of Program Mapping and Scalability with E-Kernel on Victor" Advisor: *Ma*	Manager of Modular Microsystems, IBM, Thomas J. Watson Research Center, Yorktown Heights, NY

GRAD	NAME	DISSERTATION TOPIC	FIRST POSITION
Ph.D. Recipients 1991, cont.			
Dec 91	**Tarek Sobh**	"A Discrete Event Dynamic System Model for Controlling an Observer Under Uncertainty" Advisor: *Bajcsy*	Post-doctoral research, General Robotics and Active Sensory Perception Laboratory, University of Pennsylvania, Philadelphia, PA

GRAD	NAME	DISSERTATION TOPIC	FIRST POSITION
Ph.D. Recipients 1992			
May 92	**Mario F.M. Campos**	"Robotic Exploration of Material and Kinematic Properties of Objects" Advisor: *Bajcsy*	Asst. Professor, Departmento de Ciencia da Computacao, Universidade Federal de Minas Gerais, Brazil
May 92	**Vijay Gehlot**	"Theory of Programming Languages: Concurrent Programming" Advisor: *Gunter*	Asst. Professor, Computer Science, University of Delaware, Newark, DE
May 92	**Ethel Schuster**	"Pronominal Reference to Events and Actions: Computational Foundations" Advisor: *Webber*	Returned to Israel
Aug 92	**Robert Frank**	"Formal grammar, parsing, and acquisition" Advisor: *Joshi & Kroch*	Asst. Professor, Univ. of Delaware, Newark, DE

GRAD	NAME	DISSERTATION TOPIC	FIRST POSITION
Ph.D. Recipients 1993			
May 93	**Anuj Dawar**	"Feasible Computation Through Model Theory" Advisor: *Weinstein*	Senior Research Assistant, Department of Computer Science, University of Swansea, Singleton Park, Swansea, SA2 8PP
Aug 93	**Eric Brill**	"A Corpus-Based Approach To Language Learning" Advisor: *Marcus*	Research Scientist, Spoken Language Systems Group, Laboratory for Computer Science, MIT, Cambridge MA
Aug 93	**Daniel Hardt**	"Verb Phrase Ellipsis: Form, Meaning, and Processing" Advisor: *Joshi/Webber*	Visiting Professor, Computer Science Department, Villanova University, Villanova PA
Aug 93	**Jinyoung Choi**	"The Decidability Problem For Rigid E-Unification: A New Proof And Extensions" Advisor: *Gallier*	
Aug 93	**Jiamin Zhao**	"Moving Posture Reconstruction from Perspective Projections of Jointed Figure Motion" Advisor: *Badler*	Technical Staff Member, Bell Communications Research, Morristown NJ
Dec 93	**Barbara Di Eugenio**	"Understanding Natural Language Instructions: a Computational Approach to Purpose Clauses" Advisor: *Webber*	Special Lecturer in Computational Linguistics, Computational Linguistics Program, Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA
Dec 93	**Thomas Fontaine**	"Handprinted Word Recognition Using a Hybrid of Connectionist and Procedural Methods" Advisor: *Shastri*	Research Scientist, Tudor Investment Corporation, New York, NY
Dec 93	**Joshua S. Hodas**	"Logic Programming in Intuitionistic Linear Logic: Theory, Design and Implementation" Advisor: *Miller*	Professor, Computer Science Department, Harvey Mudd College, Claremont, CA
Dec 93	**Eunyoung J. Koh**	"Automatic Synthesis of Simplified 3D Models from Detailed Data" Advisor: *Badler*	Associate, Goldman, Sachs Co., New York, NY

GRAD	NAME	DISSERTATION TOPIC	FIRST POSITION
Proposed Ph.D. Recipients Dec 1993, cont.			
Dec 93	**Michael Niv**	"A Computational Model of Syntactic Processing: Ambiguity Resolution from Interpretation" Advisor: *Steedman*	Post-doctoral candidate, Computer Science Department, Technion, Haifa, Israel
Dec 93	**Philip Resnik**	"Selection and Information: A Class-Based Approach to Lexical Relationships" Advisor: *Marcus*	Research Scientist, Sun Microsystems Laboratories, Inc., Chelmsford, MA
Dec 93	**Owen Rambow**	"Computational Models of Word Order Variation: Scrambling and Topicalization" Advisor: *Joshi/Kroch*	NATO Postdoctoral Fellow, Universite Paris 7, Paris
Dec 93	**Ron Rymon**	Diagnostic Reasoning and Planning in Exploratory-Corrective Domains Advisor: *Webber*	
Dec 93	**Marilyn Walker**	"Informational Redundancy and Resource Bounds In Dialogue" Advisor: *Joshi/Prince*	Research Scientist, Mitsubishi Cambridge Research Laboratories, Cambridge, MA

GRAD	NAME	DISSERTATION TOPIC	FIRST POSITION
Ph.D. Recipients 1994			
May 94	**Patrice Bremond-Gregoire**	"A Process Algebra of Communicating Shared Resources with Dense Time and Priorities" Advisor: *Lee?*	Director, Unix Internationalization Center, Unisys Corp., Salt Lake City, Utah
May 94	**Ron Katriel**	"Efficient Evidential Indexing of Three-Dimensional Models Using Prototypical Parts" Advisor: *?*	Senior Consultant, GARI Software Associates, Inc., Livingston, NJ

5 Report of Inventions

Tools, Applications, and Research Projects

- XTAG, a graphic interface for Tree-Adjoining Grammar development. XTAG now contains a feature-based predicative parser as well as separate morphology and syntactic databases. XTAG is distributed freely for research purposes. To obtain more information, send mail to `lex-request@linc.cis.upenn.edu`
- Jack, a 3-D interactive environment suitable for ergonomic evaluations and simulations of people with products, buildings, and situations. Jack is developed through the Center for Human Modeling and Simulation.
- TraumAID: Research and development aimed at using the tools and methods of artificial intelligence, NLP, decision theory, computer modelling and graphic simulation to improve the delivery of quality trauma care during the initial definitive management of multiple trauma.
- xv, an X11 image display program for manipulation of graphic images. To learn more about xv, contact `xv@devo.dccs.upenn.edu`
- ACL-DCI CD-ROM I, a collection of text and lexicon for natural language research, including early output from the Penn Treebank project.
- Penn Treebank CD-ROM 0.6 contains more than 4.5 million words of text, hand-annotated for parts of speech, with a 3 million word subset also hand-annotated with skeletal grammar structure.
- PEARL, an algorithm for the stochastic parsing of English.
- Tree Grep, a program that allows the user to *quickly* scan a database of tree structures for complex constructions.
- Part of Speech Taggers, one rule based, and a Markov Model based tagger.
- Fast Concordance Analysis, a tool which operates in time proportional to the pattern being searched for, rather than the size of the pattern being searched.

6 Bibliography

Please see attached Appendix A, List of Publications: October 1989 - March 1995.

7 Appendices

7.1 Appendix A: List of Publications: October 1989 - March 1995

The Convergence of Mildly Context-Sensitive Grammar Formalisms

Aravind K. Joshi, K. Vijay-Shanker, and David Weir

MS-CIS-90-01

LINC LAB 161

Investigations of classes of grammars that are *non-transformational* and at the same time highly constrained are of interest both linguistically and mathematically. Context-free grammars (CFG) obviously form such a class. CFGs are not adequate (both weakly and strongly) to characterize some aspects of language structure. Thus how much more power beyond CFG is necessary to describe these phenomena is an important question. Based on certain properties of tree adjoining grammars (TAG) an approximate characterization of class of grammars, *mildly context-sensitive grammars* (MCGS), has been proposed earlier. In this paper, we have described the relationship different grammar formalisms, all of which belong to MCSG. In particular, we have shown that head grammars (HG), combinatory categorial grammars (CCG), and linear indexed grammars (LIG) and TAG are all weakly equivalent. These formalisms are all distinct from each other at least in the following aspects: (a) the formal objects and operations in each formalism, (b) the domain of locality over which dependencies are specified, (c) the degree to which recursion and the domain of dependencies are factored, and (d) the linguistic insights that are captured in the formal objects and operations in each formalism. A deeper understanding of this convergence is obtained by comparing these formalisms at the level of the derivation structures in each formalism. We have described a formalism, the linear context-free rewriting system (LDFR), as a first attempt to capture the closeness of the derivation structures of these formalisms. LCFRs thus make the notion of MCSGs more precise. We have shown that LCFRs are equivalent to multicomponent tree adjoining grammars (MCTAGs), and also briefly discussed some variants of TAGs, lexicalized TAGs, feature structure based TAGs, and TAGs in which local domination and linear precedence are factored TAG(LD/LP).

Underestimation of Visual Texture Slant by Human Observers: A Model

M. Turner, G. Gerstein, and R. Bajcsy

MS-CIS-90-02

GRASP LAB 200

The perspective image of an obliquely inclined textured surface exhibits shape and density distortions of texture elements which allow a human observer to estimate the inclination angle of the surface. However, since the work of Gibson (1950) it has been known that, in the absence of other cues, humans tend to underestimate the slant angle of the surface, particularly when the texture is perceived as being "irregular."

The perspective distortions which affect texture elements also shift the projected spatial frequencies of the texture in systematic ways. Using a suitable local spectral filter to measure these frequency gradients, the inclination angle of the surface my be estimated. A computational model has been developed which performs this task using distributions of outputs from filters found to be a good description of simple cell receptive fields. However, for "irregular" textures the filter output distributions are more like those of regular textures at shallower angles of slant, leading the computational algorithm to underestimate the slant angle. This behavioral similarity between human and algorithm suggest the possibility that a similar visual computation is performed in cortex.

KB: A Knowledge Representation Package for Common Lisp

Jeffrey Esakov

MS-CIS-90-03

GRAPHICS LAB 30

KB is a frame-based knowledge representation package. It is written as a Common Lisp package, and is comprised of a set of functions for representing semantic knowledge and relationships among data represented. **KB** encourages the use of the object-oriented programming metaphor by requiring that a set of operators be defined for each concept (object). Inheritance is supported for both data types and for operators.

KB has a well-defined programming interface through which a user interface language can be easily developed. The semantics of **KB** are straight-forward and allow a programmer considerable flexibility in developing as application.

KB borrows heavily for the Flavors system in syntax and semantics (and in documentation).

Strength Guided Motion

Philip Lee, Susanna Wei, Jianmin Zhao, and Norman I. Badler

MS-CIS-90-04

GRAPHICS LAB 31

A methodology and algorithm is presented that generates motions imitating the way humans complete a lifting task under various loading conditions. The path taken depends on "natural" parameters: the figure geometry, the given load, the final destination, and especially, the

strength model of the agent. Additional user controllable parameters of the motion are the *comfort* of the action and the *perceived exertion* of the agent. The algorithm uses this information to incrementally compute a motion path of the end effector moving the load. It is therefore instantaneously adaptable to changing force, loading, and strength conditions. Various strategies are used to model human behavior (such as pull back, add additional joints, and jerk) that compute the driving torques as the situation changes. The strength model dictates acceptable kinematic postures. The resulting algorithm offers torque control without the tedious user expression of driving forces under a dynamics model. The algorithm runs in near-realtime and offers an agent-dependent toolkit for fast path prediction. Examples are presented for various lifting tasks, including one- and two-handed lifts, and raising the body from a seated posture.

From Simple Associations to Systematic Reasoning: A Connectionist Representation of Rules, Variables, and Dynamic Bindings

Lokendra Shastri and Venkat Ajjanagadde

MS-CIS-90-05

LINC LAB 162

Human agents draw a variety of inferences effortlessly, spontaneously, and with remarkable efficiency–as though these inferences are a *reflex* response of their cognitive apparatus. The work presented in this paper is a step toward a computational account of this remarkable reasoning ability. We describe how a connectionist system made up of simple and slow neuron-like elements can encode millions of facts and rules involving n-ary predicates and variables, and yet perform a variety of inferences within *hundreds of milliseconds*. We observe that an efficient reasoning system must represent and propagate, dynamically, a large number of variable bindings. The proposed system does so by propagating *rhythmic* patterns of activity wherein dynamic binding are represented as the *in-phase*, i.e., synchronous, firing of appropriate nodes. The mechanisms for representing and propagating dynamic bindings are biologically plausible. Neurophysiological evidence suggests that similar mechanisms may in fact be used by the brain to represent and process sensorimotor information.

Interpretive Value Analysis

David Klein

MS-CIS-90-06

This dissertation presents *interpretive value analysis (IVA)*, a framework for explaining and refining choices among competing alternatives in the context of intelligent systems. IVA increases the transparency of *value theory*, a formal model of choice, to provide a framework for modeling choices that is both formal and transparent. The components of IVA include (1) an *interpretation* of value theory that provides an intuitive yet formally sound vocabulary for talking about choices, (2) a set of *strategies for explaining choices*, and (3) a set of *strategies for refining choices*.

IVA at one addresses problems in artificial intelligence (AI) and in decision analysis (DA). From an AI perspective, IVA provides a general foundation for building formally justifiable, intelligible, modifiable systems for choosing among alternatives. A secondary contribution of the work to AI is a set of observations concerning formality and transparency; although previous approaches to modeling choices in a systems context generally have reflected a view of formality and transparency as competitive properties of representations, our experience developing IVA suggests that these properties are synergistic. Finally, the dissertation outlines a potential approach to employing other formal models in the context of intelligent systems.

From a DA perspective, IVA addresses problems of transparency. First, IVA can potentially increase the acceptance of decision-theoretic advice by providing methods for justifying that advice in intuitive terms. Second, IVA provides an approach to managing bias in parameter assessment; the framework provides users with an opportunity to observe the step-by-step effect of a parameter value on the final result, so that users' responses may potentially be influenced less by the fashion in which parameter-assessment questions are posed. Third, IVA can potentially reduce the demands on parameter-assessment methods by providing for the incremental repair of model parameters. Finally, the framework provides and approach to the problem of managing changing preferences over time.

Many of the elements of IVA are implemented in VIRTUS, a shell for building systems that choose among competing alternatives. Three practical systems have been constructed using VIRTUS, in the domains of marketing process control and medicine.

TraumAID: AI Support for the Management of Multiple Trauma

Bonnie L. Webber, John R. Clarke, Ron Rymon, Michael Niv, and María Milagros Ibáñez

MS-CIS-90-09
LINC LAB 163

This paper outlines the particular demands that multiple trauma makes on systems designed to provide appropriate decision support, and the ways that these demands are currently being met in our system, TraumAID. The demands follow from: (1) the nature of trauma and the procedures used in diagnosis, (2) the need to adjust diagnostic and therapeutic procedures to available resource levels, (3) the role of anatomy in trauma and the need for anatomical reasoning, and (4) the competing demands of multiple injuries and the consequent need for planning. We believe that these demands are not unique to multiple trauma, so that the paper may be of general interest to expert system research and development.

(This paper appears in the *Proceedings of the 1990 AAAI Symposium on Artificial Intelligence and Medicine*, Stanford University, CA March 1990).

Analyses of the Effects of Model Mismatch and Flat MMF for Estimating Particle Motion

Siu-Leong Iu

MS-CIS-90-10
GRASP LAB 204

In this report, we analyze the performance degradation due to three classes of model mismatch: parameter jumping, undermodeling and overmodeling, in estimating the particle motion by using the orthogonal polynomials to model the trajectory. We find that these model mismatches make the 'optimal estimator' to have large bias and mean squared error. For the case of undermodeling, the estimation error increases, in general, without a bound as the observation interval increases. We then propose the *Finite Lifetime Alternately Triggered Multiple Model Filter* (FLAT MMF), as a solution. FLAT MMF is a filter composed of a set of K indentical conventional state estimation filters, each triggered alternately. After the last filter is triggered, the oldest one is triggered again and so on. The structure of Multiple Model Filter is used to combine these estimates optimally, in the sense of minimum mean squared error.

We find that the ration of weightings in FLAT MMF are related to some independent non-central X^2 random variables. Consequently, we show that the FLAT MMF can provide an estimate that follows abrupt changes in the trajectory and has the small bias for undermodeling. For the case of overmodeling or the case that the trajectory model matches to the actual motion, the estimate does not degrade significantly.

A number of simulations are conducted to illustrate the estimation performance degradation due to the model mismatches for the convential Kalman filter and the performance improvement as the proposed FLAT MMF is used.

Parsing with Lexicalized Tree Adjoining Grammar

Yves Schabes
Aravind K. Joshi

MS-CIS-90-11
LINC LAB 164

Most current linguistic theories give lexical accounts of several phenomena that used to be considered purely syntactic. The information put in the lexicon is thereby increased in both amount and complexity: see, for example, lexical rules in LFG, GPSG, HPSG, Combinatory Categorial Grammars, some versions of GB theory, and Lexicon-Grammars.

We would like to take into account this fact while defining a formalism. We therefore explore the view that syntactical rules are not separated from lexical items. We say that a grammar is lexicalized if it consists of:

- A finite set of structures each associated with lexical items; each lexical item will be called the anchor of the corresponding structure; the structures define the domain of locality over which constraints are specified.

- An operation or operations for composing the structures.

There is a natural general two-step parsing strategy that can be defined for 'lexicalized' grammars. In the first stage, the parser selects a set of elementary structures associated with the lexical items in the input sentence, and in the second stage the sentence is parsed with respect to this set. The strategy is independent of the nature of the elementary structures in the underlying grammar. In principle, any parsing algorithm can be used in the second stage.

We consider Lexicalized Tree Adjoining Grammars as an instance of lexicalized grammars. We take three main types of parsing algorithms: purely top-down (as in definite clause parsing), purely bottom-up (as the Earley-type parser). For each type, we investigate if the two-step strategy provides any improvement. For the Earley-type parser, we evaluate the two-step strategy with respect

to two characteristics. First, the amount of filtering on the entire grammar is considered: once the first pass is performed, the parser uses only a subset of the grammar. Second, we evaluate the use of non-local information: The structures selected during the first pass encode the morphological value (and therefore the position in the string) of their anchor.

Interactive Real-Time Articulated Figure Manipulation using Multiple Kinematic Constraints

Cary B. Phillips, Jianmin Zhao, and Norman I. Badler

MS-CIS-90-12
GRAPHICS LAB 32

In this paper, we describe an interactive system for positioning articulated figures which uses a 3D direct manipulation technique to provide input to an inverse kinematics algorithm running in real time. The system allows the user to manipulate highly articulated figures, such as human figure models, by interactively dragging 3D "reach goals". The user may also define multiple "reach constraints" which are enforced during the manipulation. The 3D direct manipulation interface provides a good mechanism for control of the inverse kinematics algorithm and helps it to overcome problems with redundancies and singularities which occur with figures of many degrees of freedom. We use an adaptive technique for evaluating the constraints which allows us to ensure that only a certain user-controllable amount of time will be consumed by the inverse kinematics algorithm at each iteration of the manipulation process. This technique is also sensitive to the time it takes to redraw the screen, so it prevents the frame display rate of the direct manipulation from becoming too slow for interactive control.

Multi-Oriented Multi-Resolution Edge Detection

Laurent Peytavin

MS-CIS-90-13
GRASP LAB 205

In order to build an edge detector that provides information on the degree of importance spatial features represent in the visual field, I used the wavelet transform applied to two-dimensional signals and performed a multi-resolution multi-oriented edge detection. The wavelets are functions well-localized in spatial domain and in frequency domain. Thus the wavelet decomposition of a signal or an image provides outputs in which you can still extract spatial features and not only frequency components.

In order to detect edges the wavelet I chose is the first derivative of a smoothing function. I decompose the images as many times as I have directions of detection. I decided to work for the momment on the X-direction and the Y-direction only. Each step of the decomposition corresponds to a different scale. I use a discrete scale $s = 2^j$ (dyadic wavelet) and a finite number of decomposed images. Instead of scaling the filters at each step I sample the image by 2 (gain in processing time). Then, I extract the extrema, track and link them from the coarsest scale to the finest one. I build a symbolic image in which the edge-pixels are not only localized but labelled too, according to the number of appearances in the different scales and according to the contrast range of the edge. Without any arbitrary threshold I can subsequently classify the edges according to their physical properties in the scene and their degree of importance.

This process is subsequently intended to be part of more general perceptual learning procedures. The context should be: none or as little as possible a priori knowledge, and the ultimate goal is to integrate this detector in a feedback system dealing with color information, texture and smooth surfaces extraction. Then decisions must be taken on symbolic levels in order to make new interpretation or even new edge detection on ambiguous areas of the visual field.

CCSR: A Calculus for Communicating Shared Resources

Richard Gerber and Insup Lee

MS-CIS-90-16
GRASP LAB 208

The timing behavior of a real-time system depends not only on delays due to process synchronization, but also on the availability of shared resources. Most current real-time models capture delays due to process synchronization; however, they abstract out resource-specific details by assuming idealistic operating environments. On the other hand, scheduling and resource allocation algorithms used for real-time systems ignore the effect of process synchronization except for simple precedence relations between processes. To bridge the gap between these two disciplines, we have developed a formalism called Communicating Shared Resources, or CSR. This paper presents the priority-based process algebra called the Calculus for Communicating Shared Resources (CCSR), which provides an equational characterization of the CSR language. The computation model of CCSR is resource-based in that multiple resources execute synchronously, while processes assigned to the same resource are interleaved according to their priorities. CCSR possesses a prioritized

strong equivalence for terms based on strong bisimulation. The paper also describes a producer and consumer problem whose correct timing behavior depends on priority.

Animation From Instructions

Norman I. Badler, Bonnie L. Webber, Jugal Kalita, and Jeffrey Esakov

MS-CIS-90-17
GRAPHICS LAB 33
LINC LAB 165

We believe that computer animation in the form of *narrated animated simulations* can provide an engaging, effective and flexible medium for instructing agents in the performance of tasks. However, we argue that the only way to achieve the kind of flexibility needed to instruct agents of varying capabilities to perform tasks with varying demands in work places of varying layout is to drive *both* animation and narration from a *common representation* that embodies that same conceptualization of tasks and actions as *Natural Language* itself. To this end, we are exploring the use of Natural Language *instructions* to drive animated simulations. In this paper, we discuss the relationship between instructions and behavior that underlie our work and the overall structure of our system. We then describe in somewhat more detail three aspects of the system — the representation used by the *Simulator* and the *Motion Generators* used in the system.

Encoding a Dependent-Type λ-Calculus in a Logic Programming Language

Amy Felty
(INRIA Sophia-Antipolis)
Dale Miller
(University of Pennsylvania)

MS-CIS-90-18
LINC LAB 166

Various forms of typed λ-calculi have been proposed as specification languages for representing wide varieties of object logics. The *logical framework*, LF, is an example of such a dependent-type λ-calculus. A small subset of intuitionistic logic with quantification over simple type λ-calculus has also been proposed as a framework for specifying general logics. The logic of *hereditary Harrop* formulas with quantification at all non-predicate types, denoted here as hh^w, is such a meta-logic that has been implemented in both the Isabelle theorem prover and the λ-Prolog logic programming language. Both frameworks provide for specifications of logics in which details involved with free and bound variable occurrences, substitutions, eigenvaribles, and the scope of assumptions within object logics are handled correctly and elegantly at the "meta" level. In this paper, we show how LF can be encoded into hh^w in a direct and natural way by mapping the typing judgements in LF into propositions in the logic of hh^w clauses, and the proofs in one system correspond directly to proofs in the other system. Relating these two languages makes it possible to provide implementations of proof checkers and theorem provers for logics specified in LF by using standard logic programming techniques which can be used to implement hh^w.

A Contact Stress Model for Determining Forces in an Equilibrium Grasp

Pramath Raj Sinha

MS-CIS-90-19
GRASP LAB 209

Most available methods that predict the forces necessary to grasp an arbitrary object treat the object and the fingers as rigid bodies and the finger/object interface as a point contact with Coulomb friction. For statically indeterminate grasps, therefore, while it is possible to find grasps that are stable, there is no unique determination of the actual forces at the contact points and equilibrium grasps are determined as many-parameter families of solutions. Also, these models sometimes lead to phenomemologically incorrect results which, while satisfactory from a purely mathematical viewpoint, are conterintuitive and not likely to be realized in practice. The model developed here utilizes a contact-stress analysis of an arbitrarily shaped object in a multi-fingered grasp. The fingers and the object are all treated as elastic bodies and the region of contact is modeled as a deformable surface patch. The relationship between the friction and normal forces is now nonlocal and nonlinear in nature and departs from the Coulomb approximation. The nature of the constraints arising out of conditions for compatibility and static equilibrium motivated the formulation of the model as a non-linear constrained minimization problem. The total potential energy of the system is minimized, subject to the nonlinear, equality and inequality constraints on the system, using the Schittkowski algorithm. The model is able to predict the magnitude of the inwardly directed normal forces, and both the magnitude and direction of the tangential (friction) forces at each finger/object interface for grasped objects in static equilibrium. Examples in two and three dimensions are presented along with application of the model to the grasp transfer maneuver.

Extending Definite Clause Grammars with Scoping Constructs

Remo Pareschi and Dale Miller

MS-CIS-90-20

LINC LAB 167

Definite Clause Grammars (DCGs) have proved valuable to computational linguists since they can be used to specify phrase structured grammars. It is well known how to encode DCGs in Horn clauses. Some linguistic phenomena, such as *filler-gap dependencies*, are difficult to account for in a completely satisfactory way using simple phrase structured grammar. In the literature of logic grammars there have been several attempts to tackle this problem by making use of special arguments added to the DCG predicates corresponding to the grammatical symbols. In this paper we take a different line, in that we account the filler-gap dependencies by encoding DCGs within *hereditary Harrop formulas*, an extension of Horn clauses (proposed elsewhere as a foundation for logic programming) where implicational goals and universally quantified goals are permitted. Under this approach, filler-gap dependencies can be accounted for in terms of the operational semantics underlying hereditary Harrop formulas, in a way reminiscent of the treatment of such phenomena in Generalized Phrase Structure Grammar (GPSG). The main features involved in this new formulation of DCGs are mechanisms for providing scope to constants and program clauses along with a mild use of λ-terms and λ-conversion.

From Operational Semantics to Abstract Machines: Preliminary Results

John Hannan and Dale Miller

MS-CIS-90-21

LINC LAB 168

The operational semantics of functional programming languages is frequently presented using interence rules in with simple meta-logics. Such presentations of semantics can be high-level and perspicuous since meta-logics often handle numerous syntactic details in a declarative fashion. This is particularly true of the meta-logic we consider here, which includes simply types λ-terms, quantification at higher types, and β-conversion. Evaluation of functional programming languages is also often presented using low-level descriptions based on abstract machines: simple term rewriting systems in which few high-level description of evaluation using inference rules can be systematically transformed into a low-level abstract machine by removing dependencies on high-level features of the meta-logic until the resulting inference rules are so simple that they can be immediately identified as specifying an abstract machine. In particular, we present in detail the transformation of two inference rules specifying call-by-name evaluation of the untyped λ-calculus into the Krivine machine, a stack-based abstract machine that implements such evaluation. The resulting machine uses de Bruijn numerals and closures instead of formal substitution. We also comment on a construction of a simplified SECD machine implementing call-by-value evaluation. This approach to abstract machine construction provides a semantics-directed method for motivating, proving correct, and extending such abstract machines.

Estimation of General Rigid Body Motion from a Long Sequence of Images

Sui-Leong Iu

MS-CIS-90-22

GRASP LAB 210

In estimating the 3-D rigid body motion and structure from time-varying images, most of previous approaches which exploit a large number of frames assume that the rotation, and the translation in some case, are constant. For a long sequence of images, this assumption in general is not valid. In this paper, we propose a new state estimation formulation for the general motion in which the 3-D translation and rotation are modeled as the polynomials of arbitrary order. Extended Kalman filter is used to find the estimates recursively from noisy images. A number of simulations including the Monte Carlo analysis are conducted to illustrate the performance of the proposed formulation.

Structure and Intonation in Spoken Language Understanding

Mark Steedman

MS-CIS-90-23

LINC LAB 169

The structure imposed upon spoken sentences by intonation seems frequently to be orthogonal to their traditional surface-syntactic structure. However, the notion of "intonational structure" as formulated by Pierrehumbert, Selkirk, and others, can be subsumed under a rather different notion of syntactic surface structure that emerges from a theory of grammar based on a "Combinatory" extension to Categorial Grammar. Interpretations of constituents at this level are in turn directly related to "information structures", or discourse-related notions of "theme", "rheme", "focus" and presupposition". Some simplifications appear to follow for the problem of integrating syntax and other high-level modules in spoken language systems.

A Lexicalized Tree Adjoining Grammar for English

Anne Abeillé, Kathleen Bishop, Sharon Cote, and Yves Schabes

MS-CIS-90-24

LINC LAB 170

This paper presents a sizable grammar for English written in the Tree Adjoining grammar (TAG) formalism. The grammar uses a (TAG) that is both lexicalized (Schabes, Abeillé, Joshi 1988) and feature based (Vijay-Shanker, Joshi 1988). In this paper, we describe a wide range of phenomena that it covers.

A Lexicalized TAG (LTAG) is organized around a lexicon, which associates sets of elementary trees (instead of just simple categories) with the lexical items. A Lexicalized TAG consists of a finite set of trees associated with lexical items, and operations (adjunction and substitution) for composing the trees. A lexical item is called the *anchor* of its corresponding tree and directly determines both the tree's structure and its syntactic features. In particular, the trees define the domain of locality over which constraints are specified and these constraints are local with respect to their anchor. In this paper, the basic tree structures of the English LTAG are described, along with some relevant features. The interaction between the morphological and the syntactic components of the lexicon is also explained.

Next, the properties of the different tree structures are discussed. The use of S complements exclusively allows us to take full advantage of the treatment of unbounded dependencies originally presented in Joshi (1985) and Kroch and Joshi (1985). Structures for auxiliaries and raising-verbs which use adjunction trees are also discussed. We present a representation of prepositional complements that is based on extended elementary trees. This representation avoids the need for preposition incorporation in order to account for double wh-questions (preposition stranding and pied-piping) and the pseudo-passive.

A treatment of light verb constructions is also given, similar to what Abeillé (1988c) has presented. Again, neither noun nor adjective incorporation is needed to handle double passives and to account for a CNPC violations in these constructions. TAG's extended domain of locality allows us to handle, within a single level of syntactic description, phenomena that in other frameworks require either dual analyses or reanalysis.

In addition, following Abeillé and Schabes (1989), we describe how to deal with semantic non compositionality in verb-particle combinations, light verb constructions and idioms, without losing the internal syntactic composition of these structures.

The last sections discuss current work on PRO, case, anaphora and negation, and outline future work on copula constructions and small clauses, optional arguments, adverb movement and the nature of syntactic rules in a lexicalized framework.

An Improved $2n - 2$ Constant Queue Routing for an $(n \not\times n)$ Mesh

Sanguthevar Rajasekaran and Richard Overholt

MS-CIS-90-25

GRASP LAB 211

We present a $2n-2$ step routing algorithm for an $n \times n$ mesh that has a queue size of 58. The previous best known result is a routing algorithm with the same time bound but with a queue size of 672. The improvement in the queue size is possible due to (from among things) a new $3s + o(s)$ sorting algorithm for an $s \times s$ mesh.

An Optimal Randomized Algorithm for Selection on the Hypercube

Sanguthevar Rajasekaran

MS-CIS-90-26

GRASP LAB 212

We show that selection on an input of size N can be performed on a P-node hypercube in time $O(N/P)$, provided each node can process all the incident edges in one unit of time (this model is called the *parallel model* and has been assumed by previous researchers (e.g., [15])). This result is important in view of a lower bound of Plaxton that implies selection takes $\Omega((N/P)\log\log P + \log P)$ time on a P-node hypercube if each node can process only one edge at a time (this model is referred to as the *sequential model*).

Minimax Estimation of a Discrete Location Parameter for a Continuous Distribution *(Dissertation)*

Raymond A. McKendall

MS-CIS-90-28

GRASP LAB 214

The subject of this research is the following stochastic model: $Z = \theta + V$. The random variable Z is a measurement of the discrete location parameter θ in continuous, additive noise V. The possible values of θ are $0, \pm u, \pm 2u, \ldots, \pm Nu$, where N is a positive integer and u is a positive number. The distribution of Z, denoted $F_Z(\cdot \mid \theta)$, is continuous and increasing on $\Re$ and has a continuous density. In a standard-estimation problem, the distribution $F_Z(\cdot \mid \theta)$ is known. In a robust-estimation problem, the distribution $F_Z(\cdot \mid \theta)$ is uncertain: It is

an unknown member of an uncertainty class of distributions. In either case, the goal of this research is to find a minimax estimator of the location parameter θ from an observation Z. This goal subsumes the identification of equalizer decision rules and Bayes decision rules for estimating θ. The loss function in the underlying statistical decision problem is the zero-one loss function with error tolerance e (a non-negative integer):

$$L_e(\theta, \hat{\theta}) := \begin{cases} 0 \text{ if } |\theta - \hat{\theta}| \leq eu \\ 1 \text{ if } |\theta - \hat{\theta}| > eu \end{cases}$$

This loss is independent of the noise distribution.

Different additional assumptions define the specific problems considered. Chapters 2, 3, and 4 consider standard estimation. Chapter 2 discusses standard estimation as a problem in statistical decision theory. Chapter 3 assumes that the observable Z has a monotone likelihood ratio. Its decision rules are monotonic. In contrast, chapter 4 assumes that the observable has a Cauchy distribution, which does not have a monotone likelihood ratio. Its decision rules are not monotonic. Chapters 5, 6, and 7 consider robust estimation. Chapter 5 formulates robust estimation as a problem in statistical decision theory. Chapter 6 extends the results for the standard estimation of Chapter 3 to robust estimation, and consequently its decision rules are monotonic. Chapter 7 considers a simple robust-estimation problem for which the extension of chapter 6 does not apply. Its decision rules are not monotonic. In all of these problems, the decision rules are non-randomized step functions. The steps occur at points determined by nonlinear systems of equations in the noise distribution.

Chapters 2 - 7 summarize the main results. Chapters 8 - 21 give the analysis.

Computation and Linguistics Theory: A Government Binding Theory Parser using Tree Adjoining Grammar (Dissertation Proposal)

Robert Frank

MS-CIS-90-29

LINC LAB 171

Government Binding (GB) theory, as a competence theory of grammar, is intended to define what a speaker's knowledge of language consists of. The theory proposes a system of innate principles and constraints which determine the class of possible languages and once instantiated by the parameter values for a given language, the class of well-formed sentences of that language [Chomsky, 1981].

In this thesis, I address the problem of how this knowledge of language is put to use. The answer I give to this question takes the shape of an implemented computational model, a parser, which utilizes the formulation of knowledge of language as proposed in GB theory. GB as a theory of grammar poses a particular problem for instantiation within a cognitively feasible computational model. It has a rich deductive structure whose obvious direct implementation as a set of axioms in a first order theorem prover runs up against the problem of undecidability. Thus, if we accept GB theory as psychologically real, and thus as functioning casually with respect to linguistic processing, there seems to be a paradox: we need a way of putting our knowledge of language, represented in GB theory, to use in a processing theory in an efficient manner.

I will suggest a way out of this paradox. I propose to constrain the class of possible grammatical principles by requiring them to be statable over a linguistically and mathematically motivated domain, that of a tree adjoining grammar (TAG) elementary tree. The parsing process consists of the construction of such primitive structures, using a generalization of licensing relations of proposed in [Abney, 1986], and checking that the constraints are satisfied over these local domains. Since these domains are of bounded size, the constraints will be checkable in constant time and we will be guaranteed efficient, linear time, parsing. Additionally, the incrementality of the construction of the TAG elementary trees is consistent with intuitions of incremental semantic interpretation.

Segmentation as the Search for the Best Description of the Image in Terms of Primitives

Aleš Leonardis, Alok Gupta, and Ruzena Bajcsy

MS-CIS-90-30

GRASP LAB 215

Segmentation of images has long been considered in computer vision as an important but extremely difficult problem In this paper we present a new paradigm for the segmentation of images into piecewise continuous patches. Data aggregation is performed via model recovery in terms of variable-order bi-variate polynomials using iterative regression. All the recovered models are potential candidates for the final description of the data. Selection of the models is achieved through a maximazation of quadratic Boolean problem. The procedure can be adapted to prefer certain kind of description (one which describes more data points, or has smaller error, or has lower order model). We have developed a fast optimization procedure for model selection. The major novelty of the approach is in combining model extraction and model selection in a dynamic way. Partial recovery of the mod-

els is followed by the optimization (selection) procedure where only the "best" models are allowed to develop further. The results obtained in this way are comparable with the result obtained when using the selection module only after all the models are fully recovered, while the computational complexity if significantly reduced. We test the procedure on real range and intensity images.

Representing Objects in a Logic Programming Language with Scoping Constructs

Joshua S. Hodas and Dale Miller

MS-CIS-90-31

LINC LAB 172

We present logic programming language that uses implications and universal quantifiers in goals and in the bodies of clauses to provide a simple scoping mechanism for program clauses and constants. Within this language it is possible to define a simple notion of parametric module and local constant. Given this ability to structure programs, we explore how object-oriented programming, where objects are viewed as abstractions with behaviors, state, and inheritance, might be accommodated. To capture the notion of mutable state, we depart from the pure logic setting by adding a declaration that certain local predicates are deterministic (they succeed at most once). This declaration, along with a goal-continuation passing style of programming is adequate to model the state of objects. We also examine a few aspects of how having objects embedded in logic programming can be used to enrich the notion of object: for example, objects may be partial (that is, may contain free variables) and nondeterministic, and it is possible not only to search for objects with certain properties but also to do hypothetical reasoning about them.

A Framework for Observing a Manipulation Process

Ruzena Bajcsy and Tarek Sobh

MS-CIS-90-34

GRASP LAB 216

We propose a system for observing a robot hand manipulating an object. A discrete event dynamic system is used as a model for the mainpulation process. A framework for the hand/object relationship developed and a stabilizing observer is constructed for the system. We describe low-level modules for recognizing the "events" that causes state transitions within the dynamic manipulation system. Our system uses different tracking mechanisms in order to control the observation process in an efficient and stable manner.

Polymorphic Rewriting Conserves Algebraic Strong Normalization

Val Breazu-Tannen and Jean Gallier

MS-CIS-90-36

LOGIC & COMPUTATION 19

We study combinations of many-sorted algebraic term rewriting systems and polymorphic lambda term rewriting. Algebraic and lambda terms are mixed by adding the symbols of the algebraic signature to the polymorphic lambda calculus, as higher-order constants.

We show that if a many-sorted algebraic rewrite system R is strongly normalizing (terminating, noetherian), then $R + \beta + \eta +$ type-β + type-η rewriting of mixed terms is also strongly normalizing. The result is obtained using a technique which generalizes Girard's "candidats de reductibilité", introduced in the original proof of strong normalization for the polymorphic lambda calculus.

Polymorphic Rewriting Conserves Algebraic Confluence

Val Breazu-Tannen and Jean Gallier

MS-CIS-90-37

LOGIC & COMPUTATION 20

We study combinations of many-sorted algebraic term rewriting systems and polymorphic lambda term rewriting. Algebraic and lambda terms are mixed by adding the symbols of the algebraic signature to the polymorphic lambda calculus, as higher-order constants.

We show that if a many-sorted algebraic rewrite system R has the Church-Rosser property (is confluent), then $R + \beta +$ type-β + type-η rewriting of mixed terms has the Church-Rosser property too.

η reduction does not commute with algebraic reduction, in general. However, using long normal forms, we show that if R is canonical (confluent and strongly normalizing) then equational provability from $R + \beta + \eta +$ type-β + type-η is still decidable.

How to Decide from the First View Where to Look Next

Jasna Maver
Ruzena Bajcsy

MS-CIS-90-39
GRASP LAB 219

The task we want to achieve is the description of a random arrangement of unknown objects in a scene. First, a complete spatial map of the scene has to be acquired. To resolve the ambiguities that are caused by occlusions in range images, we need to take sensor measurements from several different views. We have limited ourselves to the images obtained by a laser scanning system which possesses certain features — occluded regions which are easily detected and can be used in designing an efficient algorithm. We develop a strategy to determine the sequence of different views using the information in a narrow zone around the occluded regions. Occluded regions are approximated by polygons. Based on the height information of the border of the occluded regions and geometry of the edges of the polygonal approximation, the next views are determined.

Structure and Intonation

Mark Steedman

MS-CIS-90-45
LINC LAB 174

Rules for assigning phrasal intonation to sentences are often assumed to require an autonomous level of "intonational structure", distinct from what is usually thought of as surface syntactic structure. The present paper argues that the requisite notion of structure can be subsumed under the generalized notion of surface structure that emerges from the combinatory extension of Categorial Grammar. According to this theory, the syntactic structures and the intonational structures of English are one, and can be captured in a single unified grammar. The interpretations that the grammar provides for such constituents correspond to the entities and open propositions that are concerned in certain discourse-related aspects of intonational meaning that have variously been described as "theme" and "rheme", "given" and "new", or "presupposition" and "focus".

Mathematical and Computational Aspects of Lexicalized Grammars (Dissertation)

Yves Schabes

MS-CIS-90-48
LINC LAB 179

Most current linguistic theories give lexical accounts of several phenomena that used to be considered purely syntactic. The information put in the lexicon is thereby increased both in amount and complexity. We explore the view that syntactic rules are not separated from lexical items. In this approach, each elementary structure is associated with a lexical item called the anchor. These structures specify extended domains of locality (as compared to context-free grammars) over which constraints can be stated. The 'grammar' consists of a lexicon where each lexical item is associated with a finite number of structures for which that item is the anchor. There are 'rules' which tell us how these structures are composed. A grammar of this form will be said to be lexicalized.

The process of lexicalization of context-free grammars (CFGs) constrained by linguistic requirements forces us to use operations for combining structures that make the formalism fall in the class of mildly context sensitive languages. We show that substitution, the combining operation corresponding to CFGs, does not allow one to lexicalize CFGs but the combination of substitution and adjunction does. We show how tree-adjoining grammar (TAG) is derived from the lexicalization process of CFGs. Then we show that TAGs are closed under lexicalization and we illustrate the main structures found in a lexicalized TAG for English. The properties of TAGs permit us to encapsulate diverse syntactic phenomena in a very natural way. TAG's extended domain of locality and its factoring of recursion from local dependencies enable us to localize many syntactic dependencies (such as filler-gap) as well as semantic dependencies (such as predicate-arguments).

We investigate the processing of lexicalized TAGs. We first present two general practical parsers that follow Earley-style parsing. They are practical parsers for TAGs because, as for CFGs, the average behavior of Earley-type parsers is superior to its worst case complexity. They are both left to right bottom-up parsers that use top-down predictions but they differ in the way the top down prediction is used.

Then we explain the building of a set of deterministic bottom-up left to right parsers which analyze a subset of tree-adjoining languages. The LR parsing strategy for CFGs is extended to TAG by using a machine, called Bottom-up Embedded Push Down Automaton (BEPDA), that recognizes in a bottom-up fashion the set of tree-adjoining languages (and exactly this set).

Finally we show how lexicalized grammars suggest a natural two-step parsing strategy. We consider lexicalized TAGs as an instance of lexicalized grammar and we examine the effect of the two-step parsing strategy on main types of parsing algorithms.

TraumAID: Reasoning and Planning in the Initial Definitive Management of Multiple Injuries

Bonnie L. Webber
John R. Clarke
Michael Niv
Ron Rymon
María Milagros Ibáñez

MS-CIS-90-50
LINC LAB 180

The TraumAID system has been designed to provide computerized decision support to optimize the *initial definitive management* of acutely injured patients *after resusitation and stabilization*. The currently deployed system, TraumAID 1.0, addresses penetrating injuries to the abdomen and to the

chest. Our experience with TraumAID 1.0 has demonstrated some major deficiencies in rule-based reasoners faced with problems of both diagnosis and treatment. To address these deficiencies, we have redesigned the system (TraumAID 2.0), factoring it into two modules: (1) a *rule-based reasoner* embodying the knowledge and logical machinery needed to link clinical evidence to diagnostic and therapeutic *goals* and (2) a *planner* embodying the global knowledge and logical machinery needed to create a plan that address *combinations* of them. After describing TraumAID 2.0, we discuss an extension of the traumAID interface(*critique mode interaction*) that may improve its acceptability in a clinical setting. We close with a brief discussion of management support in resource-limited environments, which is an important issue in the time-critical context of multiple trauma.

Image Understanding at the GRASP Laboratory

Ruzena Bajcsy

MS-CIS-90-52
GRASP LAB 225

Research in the GRASP Laboratory has two main themes, parameterized multi-dimensional segmentation and robust decision making under uncertainty. The multi-dimensional approach interweaves segmentation with representation. The data is explained as a best fit in view of parametric primitives. These primitives are based on physical and geometric properties of objects and are limited in number. We use primitives at the volumetric level, the surface level, and the occluding contour level, and combine the results. The robust decision making allows us to combine data from multiple sensors. Sensor measurements have bounds based on the physical limitations of the sensors. We use this information without making a priori assumptions of distributions within the intervals of a priori assumptions of the probability of a given result.

A Logic Programming Language with Lambda-Abstraction, Function Variables, and Simple Unification

Dale Miller

MS-CIS-90-54
LINC LAB 182

It has been argued elsewhere that a logic programming language with function variables and lambda-abstractions within terms makes a very good meta-programming language, especially when an object language contains notions of bound variables and scope. The lambda Prolog logic programming language and the closely related Elf and Isabelle systems provide meta-programs with both function variables and lambda-abstractions by containing implementations of higher-order unification. In this paper, we present a logic programming language, called L-lambda, that also contains both function variables and lambda-abstractions, but certain restriction are placed on occurrences of function variables. As a result, an implementation of L-lambda does not need to implement full higher-order unification. Instead, an extension to first-order unification that respects bound variable names and scopes is all that is required. Such unification problems are shown to be decidable and to possess most general unifiers when unifiers exist. A unification algorithm and logic programming interpreter are described and proved correct. Several examples of using L-lambda as a meta-programming language are presented.

[[To appear in *Extensions of Logic Programming* edited by Peter Schroeder-Heister, Lecture Notes in Artificial Intelligence, Springer-Verlag. Supported in part by grants ONR N00014-88-K-0633, NSF CCR-87-05596, and DARPA N00014-85-K-0018.]]

A Logic Programming Language with Lambda-Abstraction, Function Variables, and Simple Unification

Dale Miller

MS-CIS-90-54
LINC LAB 182

It has been argued elsewhere that a logic programming language with function variables and lambda-abstractions within terms makes a very good meta-programming language, especially when an object language contains notions of bound variables and scope. The lambda Prolog logic programming language and the closely related Elf and Isabelle systems provide meta-programs with both function variables and lambda-abstractions by containing implementations of higher-order unification. In this paper, we present a logic programming language, called L-lambda, that also contains both function variables and lambda-abstractions, but certain restriction are placed on occurrences of function variables. As a result, an implementation of L-lambda does not need to implement full higher-order unification. Instead, an extension to first-order unification that respects bound variable names and scopes is all that is required. Such unification problems are shown to be decidable and to possess most general unifiers when unifiers exist. A unification algorithm and logic programming interpreter are described and proved correct. Several examples of using L-lambda as a meta-programming language are presented.

[[To appear in *Extensions of Logic Programming* edited by Peter Schroeder-Heister, Lecture Notes in Artificial Intelligence, Springer-Verlag. Supported in part by grants ONR N00014-88-K-0633, NSF CCR-87-05596, and DARPA N00014-85-K-0018.]]

Non-Monotonic Decision Rules for Sensor Fusion

Raymond McKendall
Max Mintz

MS-CIS-90-56
GRASP LAB 228

This article describes non-monotonic estimators of a location parameter θ from a noisy measurement $Z = \theta + V$ when

the possibly values of θ have the form $0, \pm 1, \pm 2 \cdots, \pm n$. If the noise V is Cauchy, then the estimator is a non-monotonic step function. The shape of this rule reflects the non-monotonic shape of the likelihood ratio of a Cauchy random variable. If the noise V is Gaussian with one of two possible scales, then the estimator is also a non-monotonic shape of the likelihood ratio of the marginal distribution of Z given θ under a least-favorable prior distribution.

Structure and Ostension in the Interpretation of Discourse Deixis

Bonnie Lynn Webber

MS-CIS-90-58
LINC LAB 183

This paper examines demonstrative pronouns used as *deictics* to refer to the interpretation of one or more clauses. Although this usage is frowned upon in style manuals such as (Strunk and White, 1959), who say

"*This.* The pronoun *this*, referring to the complete sense of a preceding sentence or clause, cannot always carry the load and so may produce an imprecise statement."

it is nevertheless very common in written text. Handling this usage poses a problem for Natural Language Understanding systems. The solution I propose is based on distinguishing between what can be *pointed to* and what can be *referred to* by virtue of pointing. I argue that a restricted set of *discourse segments* yield what such demonstrative pronouns can point to in the *discourse model* and a restricted set of what Nunberg (1979) has called *referring functions* yield what they can refer to by virtue of that pointing.

An Extension to ML to Handle Bound Variables in Data Structures: (Preliminary Report)

Dale Miller

MS-CIS-90-59
LINC LAB 184

Most conventional programming languages have direct methods for representing first-order terms (say, via concrete datatypes in ML). If it is necessary to represent structures containing bound variables, such as lambda-terms, formulas, types, or proofs, these must first be mapped into first-order terms, and then a significant number of auxiliary procedures must be implemented to manage bound variable names, check for free occurrences, do substitution, test for equality modulo alpha-conversion, etc. We shall show how the applicative core of the ML programming language can be enhanced so that lambda-terms can be represented more directly and so that the enhanced language, called ML-lambda, provides a more elegant method of manipulating bound variables within data structures. In fact, the names of bound variables will not be accessible to the ML-lambda programmer. This extension to ML involves the following: introduction of the new type constructor $'a \Rightarrow' b$ for the type of lambda-terms formed by abstracting a parameter of type $'a$ out of a term of type $'b$; a very restricted and simple form of higher-order pattern matching; a method for extending a given data structure with a new constructor; and, a method for extending function definitions to handle such new constructors. We present several examples of ML-lambda programs.

[[Appears in the Proceedings of the Logical Frameworks BRA Workshop, Nice, June 1990.]]

Using Kinds to Represent Heterogeneous Collections in a Static Type System (Extended Abstract)

Peter Buneman
(University of Pennsylvania)
Atsushi Ohori
(University of Glasgow)

MS-CIS-90-62
LOGIC & COMPUTATION 22

We consider the problem of representing heterogeneous collections of objects in a type polymorphic programming language in such a way that common properties of members of a collection, such as having commonly named filed with a common type can be expressed in the type system. The use of such collections is widespread in object-oriented and database programming and has so far been achieved in statically typed systems only through the use of a single *dynamic* type, which effectively hides all the structure of a value. In this paper we exploit a system of types and *kinds* (sets of types) to represent dynamic values with some known properties. The type system is shown to be sound and to have a complete type inference algorithm.

Polymorphism and Type Inference in Database Programming

Peter Buneman
(University of Pennsylvania)
Atsushi Ohori
(University of Glasgow)

MS-CIS-90-64
LOGIC & COMPUTATION 23

The polymorphic type system of ML can be extended in two ways that make it appropriate as the basis of a database programming language. The first is an extension to the language of types that captures the polymorphic nature of field selection; the second is a technique that generalizes relational operators to arbitrary data structures. The combination provides a statically typed language in which relational databases may be cleanly represented as typed structures. As in ML types are inferred, which relieves the programmer of making the rather complicated type assertions that may be required to express the most general type of a program that involves field selection and generalized relational operators.

It is also possible to use these ideas to implement various aspects of object-oriented databases. By implementing

database objects as reference types and generating the appropriate views – sets of structures with "identity" – we can achieve a degree of static type checking for object-oriented databases. Moreover it is possible to exploit the type system to check the consistency of object-oriented classes (abstract data types with inheritance). A prototype language based on these ideas has been implemented. While it lacks some important practical features, it demonstrates that a wide variety of database structures can be cleanly represented in a polymorphic programming language.

Human Factors Simulation Research at the University of Pennsylvania

Norman I. Badler
MS-CIS-90-67
GRAPHICS LAB 34

Jack is a Silicon Graphics Iris 4D workstations-based system for the definition, manipulation, animations, and human factors performance analysis of simulated human figures. Built on a powerful representation for articulated figures, *Jack* offers the interactive user a simple, intuitive, and yet extremely capable interface into any 3-D articulated world. *Jack* incorporates sophisticated systems for anthropometric human figure generation, multiple limb positioning under constraints, view assessment, and strength model-based performance simulation of human figures. Geometric workplace models may be easily imported into *Jack*. Various body geometries may be used, from simple polyhedral volumes to contour-scanned real figures. High quality graphics of environments and clothed figures are easily obtained. Descriptions of some work in progress are also included.

Sensor-Fusion with Statistical Decision Theory: A Prospectus of Research in the GRASP Lab

Raymond McKendall
Max Mintz
MS-CIS-90-68
GRASP LAB 234

The purpose of this report is to describe research in sensor fusion with statistical decision theory in the GRASP Lab, Department of Computer and Information Science, University of Pennsylvania. This report is thus a tutorial overview of the general research problem, the mathematical framework for the analysis, the results of specific research problems, and directions of future research. The intended audience for this report includes readers seeking a self-contained summary of the research as well as students considering study in this area. The prerequisite for understanding this report is familiarity with basic mathematical statistics.

Coordination of Two-Arm Pushing

Xiaoping Yun
MS-CIS-90-71
GRASP LAB 237

Coordination of two manipulators performing the task of transporting objects is studied in this paper. Each manipulator is equipped with end effector — a flat surface palm. Grasping is achieved by the two palms pushing an object from two ends. The task requires simultaneous control of the object motion and the interaction force. The control of the interaction force is needed to ensure that the object is not dropped and to avoid excessive pressing. The motion and force control problem is further complicated by the presence of unilateral constraints since the manipulators can only push the object. This paper describes a control method which utilizes a state feedback to decouple position control and force control loops. A force control planning algorithm is also proposed which ensures the satisfaction of unilateral constraints. The effectiveness of the control method is verified by simulations.

The Mixed Powerdomain

Carl A. Gunter
MS-CIS-90-75
LOGIC & COMPUTATION 25

This paper introduces an operator M called the *mixed powerdomain* which generalizes the convex (Plotkin) powerdomain. The construction is based on the idea of representing partial information about a set of data items using a PAIR of sets, one representing partial information in the manner of the upper (Smyth) powerdomain and the other in the manner of the lower (Hoare) powerdomain where the components of such pairs are required to satisfy a consistency condition. This provides a richer family of meaningful partial descriptions than are available in the convex powerdomain and also makes it possible to include the *empty set* in a satisfactory way. The new construct is given a rigorous mathematical treatment like that which has been applied to the known powerdomains. It is proved that M is a continuous functor on bifinite domains which is left adjoint to the forgetful functor from a category of continuous structures called MIX ALGEBRAS. For a domain D with a coherent Scott topology, elements of M(D) can be represented as pairs (U,V) where U is a compact upper subset of D and V is a closed subset of D and the downward closure of U intersect V is equal to V. A Stone dual characterization of M is also provided.

Progressive Horizon Planning

Ron Rymon and Bonnie L. Webber
(University of Pennsylvania)
John R. Clarke
(Medical College of Pennsylvania)
MS-CIS-90-76
LINC LAB 185

In an earlier paper [Rymon, 1989] we showed how domain localities and regularities can be used to reduce the complexity of finding a trauma management plan that satisfies a *set* of diagnostic and therapeutic goals. Here, we present another planning idea – *Progressive Horizon* – useful for optimizing such

plans in domains where planning can be regarded as an incremental process, continuously interleaved with situation-goals analysis and plan execution. In such domains, planned action cannot be delayed until all essential information is available: A plan must include actions intended to gather information as well as ones intended to change the state of the world.

Interleaving planning with reasoning and execution, a progressive horizon planner constructs a plan that answers *all currently known* needs but has only its first few actions optimized (those within its planning horizon). As the executor carries out actions and reports back to the system, the current goals and the plan are updated based on actual performance and newly discovered goals and information. The new plan is then optimized within a newly set horizon.

In this paper, we describe those features of a domain that are salient for the use of a progressive horizon planning paradigm. Since we believe that the paradigm may be useful in other domains, we abstract from the exact techniques used by our program to discuss the merits of the general approach.

Parallel Kalman Filtering on the Connection Machine

Michael A. Palis
(University of Pennsylvania)
Donald K. Krecker
(General Electric Companies)

MS-CIS-90-81
LINC LAB 186

A parallel algorithm for square root Kalman filtering is developed and implemented on the Connection Machine (CM). The algorithm makes efficient use of *parallel prefix* or *scan* operations which are primitive instructions in the CM. Performance measurements show that the CM filter runs in time linear in the state vector size. This represents a great improvement over serial implementations which run in cubic time. A specific multiple target tracking application is also considered, in which several targets (e.g., satellites, aircrafts and missiles) are to be traced simultaneously, each requiring one or more filters. A parallel algorithm is developed which, for fixed size filters, runs in constant time, independent of the number of filters simultaneously processed.

Description Succinctness of Some Grammatical Formalisms for Natural Language

Michael A. Palis
(University of Pennsylvania)
Sunil Shende
(University of Nebraska)

MS-CIS-90-82
LINC LAB 187

We investigate the problem of describing languages compactly in different grammatical formalisms for natural languages. In particular, the problem is studied from the point of view of some newly developed natural language formalisms like linear control grammars (LCGs) and tree adjoining grammars (TAGs); these formalisms not only generate non-context-free languages that capture a wide variety of syntactic phenomena found in natural language, but also have computationally efficient polynomial time recognition algorithms. We prove that the formalisms enjoy the property of unbounded succintness over the family of context-grammars, i.e. they are, in general, able of provide more compact representations of natural languages as compared to standard context-free grammars.

Issues in Facial Animation

Catherine Pelachaud
Norman I. Badler
Mark Steedman

MS-CIS-90-88
GRAPHICS LAB 36

Our goal is to build a system of 3-D animation of facial expressions of emotion correlated with the intonation of the voice. Up till now, the existing systems did not take into account the link between these two features. Many linguists and psychologists have noted the importance of spoken intonation for conveying different emotions associated with speakers' messages. Moreover, some psychologists have found some universal facial expressions linked to emotions and attitudes. We will look at the rules that control these relations *(intonation/emotions and facial expressions/emotions)* as well as the coordination of these various modes of expressions. Given an utterance, we consider how the message (what is *new/old* information in the given context) transmitted through the choice of accents and their placement, are conveyed through the face. The facial model integrates the action of each muscle or group of muscles as well as the propagation of the muscles' movement. It is also adapted to the FACS notation (Facial Action Coding System) created by P. Ekman and W. Friesen to describe facial expressions. Our first step will be to enumerate and to differentiate facial movements linked to emotions from the ones linked to conversation. Then, we will examine what the rules are that drive them and how their different actions interact.

Key words: facial animation, emotion, intonation, coarticulation, conversational signals

On the Convergence Time of Simulated Annealing

Sanguthevar Rajasekaran

MS-CIS-90-89
GRASP LAB 242

Simulated Annealing is a family of randomized algorithms used to solve many combinatorial optimization problems. In practice they have been applied to solve some presumably hard (e.g., NP-complete) problems. The level of performance obtained has been promised [5, 2, 6, 14]. The success of its heuristic technique has motivated analysis of this algorithm from a theoretical point of view. In particularly, people have looked at the convergence of this algorithm. They have show

(see e.g., [10]) that this algorithm converges in the limit to a globally optimal solution with probability 1. However few of these convergence results specify a time limit within which the algorithm is guaranteed to converge(with some high probability, say). We present, for the first time, a simple analysis of SA that will provide a time bound for convergence with overwhelming probability. The analysis will hold no matter what annealing schedule is used. Convergence of Simulated Annealing in the limit will follow as a corollary to our time convergence proof.

In this paper we also look at optimization problems for which the cost function has some special properties. We prove that for these problems the convergence is much faster. In particular, we give a simpler and more general proof of convergence for Nested Annealing, a heuristic algorithm developed in [12]. Nested Annealing is based on defining a graph corresponding to the given optimization problem. If this graph is 'small separable,' they [12] show that Nested Annealing will converge 'faster'.

For arbitrary optimization problem, we may not have any knowledge about the 'separability' of its graph. In this paper we give tight bounds for the 'separability' of a random graph. We then use these bounds to analyze the expected behavior of Nexted Annealing on an arbitrary optimization problem. The 'separability' bounds we derive in this paper are of independent interest and have the potential of finding other applications.

CLiFF Notes #1: Research in Natural Language Processing at the University of Pennsylvania Biannual Report: Fall 1990
Contributors: Students & Faculty
Editor: Elizabeth Levison
MS-CIS-90-95
LINC LAB 190

CLiFF is the Computational Linguists' Feedback Forum. This technical report is a collection of short abstracts by students and faculty in which they describe their work currently in progress. These presentations of the work being done in Natural Language Processing at the University of Pennsylvania should provide insight into the diversity of work at Penn, and the strong ties between the departments of Computer Science, Psychology and Linguistics.

Observing A Moving Agent

Ruzena Bajcsy
Tarek Sobh
MS-CIS-91-01
GRASP LAB 247

We address the problem of observing a moving agent. In particular, we propose a system for observing a manipulation process, where a robot hand manipulates an object. A discrete event dynamic system (DEDS) from work is developed for the hand-object interaction over time and a stabilizing observer is constructed. Low-level modules are developed for recognizing the "events" that causes state transitions within the dynamic manipulation system. The work examines closely the possibilities for errors, mistakes and uncertainties in the manipulation system, observer construction process and event identification mechanisms. The system utilizes different tracking techniques in order to observe the task in an *active, adaptive* and *goal-directed* manner.

Model Based Teleoperation To Eliminate Feedback Delay NSF Grant BCS89-01352 First Report

Richard P. Paul
Janez Funda
Simeon Thierry
Thomas Lindsay
Masahiko Hashimoto
MS-CIS-91-02
GRASP LAB 248

We are conducting research in the area of teleoperation with feedback delay. Delay occurs with earth-based teleoperation in space and with surface-based teleoperation with untethered submersibles when acoustic communication links are involved. the delay in obtaining position and force feedback from remote slave arms makes teleoperation extremely difficult. We are proposing a novel combination of graphics and manipulator programming to solve the problem by interfacing a teleoperator master arm to a graphics based simulator of the remote environment coupled with a robot manipulator at the remote, delayed site. the operator's actions will be monitored to provide both kinesthetic and visual feedback and to generate symbolic motion commands to the remote slave. the slave robot will then execute these symbolic commands delayed in time. While much of a task will proceed error free, when an error does occur the slave system will transmit data back to the master and the master environment will be "reset" to the error state.

Interactive Behaviors For Bipedal Articulated Figures

Cary B. Phillips
Norman I. Badler
MS-CIS-91-03
GRAPHICS LAB 37

We describe techniques for interactively controlling bipedal articulated figures through kinematic constraints. These constraints model certain behavioral tendencies which capture some of the characteristics of human-like movement, and give us control over such elements as the figures' balance and stability. They operate in near real-time, so provide behavioral control for interactive manipulation. These constraints form the basis of an interactive motion-generation system that allows the active movement elements to be layered on top of the passive behavioral constraints.

Hands: Human To Robotic

Sanjay Agrawal
MS-CIS-91-04
GRASP LAB 249

Hands have for centuries been recognized as a fundamental tool for humans to gain an understanding of their environment and at the same time be able to manipulate it. In this presentation we will look at various studies made on the functionality and use of the human hand and examine the different approaches to analyzing and classifying human grasps and building a taxonomy of these grasps. We study the anatomy of the human hand, and examine experiments performed to understand the how gripping forces are applied when lifting objects, and the methods extraction of haptic information, by humans.

We discuss issues involved in the building of electromechanical manipulators and some of the mathematics used in analyzing the suitability of a design. We look at one of the earliest designs of a computer controlled articulated gripper, as well as two of the most prevalent designs in today's research world, the Stanford/JPL hand and the Utah/MIT had. Finally, we show why a more fundamental understanding of how human grasping works will help us design more useful manipulators.

An Hand-Eye Arm Coordinated System

Sanjay Agrawal
Ruzena Bajcsy
Vijay Kumar
MS-CIS-91-05
GRASP LAB 250

In this paper we present the description and experiments with a tightly coupled Hand-Eye-Arm manipulatory system. We explain the philosophy and the motivation for building a tightly coupled system that actually consists of very autonomous modules that communicate with each other via a central coordinator. We describe each of the modules in the system and their interactions with each other. We highlight the need for sensory driven manipulation, and explain how the above system, where the hand is equipped with multiple tactile sensors, is capable of both manipulating unknown objects, but also detecting and complying in the case of collisions. We explain the partition of the control of the system into various closed loops, representing coordination both at the level of gross manipulator motions as well as fine motions. We describe the various modes that the system can work in , as well as some of the experiments that are being currently performed using this system.

Emulation Of A PRAM On Leveled Networks
David S. L. Wei
MS-CIS-91-06
GRASP LAB 251

We present efficient emulations of the CRCW PRAM on a large class of processor interconnection networks called *leveled networks*. This class includes the *star graph* and the n-way shuffle, which have the interesting property that the network diameter is *sub-logarithmic* in the network size. We show that a CRCW PRAM can be emulated optimally on these networks (i.e., each emulation step takes time linear in the network diameter). this is the first result that demonstrates PRAM emulation in less than logarithmic time.

We also present an efficient emulation of the CRCW PRAM on an nxn mesh. Although an $O(n)$-time emulation algorithm for the mesh is known, the underlying constant in the run-time is large, making it impractical. We give an improved emulation algorithm whose time bound is only $4n + o(n)$.

Symbolic Simulator/Debugger For The Systolic/Cellular Array Processor
Janez Funda
MS-CIS-91-07
GRASP LAB 252

Descriptive Complexity Approaches To Inductive Inference
Kevin Atteson
MS-CIS-91-08
GRASP LAB 253

We present a critical review of descriptive complexity approaches to inductive inference. Inductive inference is defined as any process by which a model of the world is formed from observations. The descriptive complexity approach is a formalization of Occam's razor: choose the simplest model consistent with the data. Descriptive complexity as defined by Kolmogorov, Chaitin and Solomonoff is presented as a generalization of Shannon's entropy. We discuss its relationship with randomness and present examples. However, a major result of the theory is negative: descriptive complexity is uncomputable.

Rissanen's minimum description length (MDL) principle is presented as a restricted form of the descriptive complexity which avoids the uncomputability problem. We demonstrate the effectiveness of MDL through its application to AR processes. Lastly, we present and discuss LeClerc's application of MDL to the problem of image segmentation.

Investigating A Proof-Theoretic Meta-Language For Functional Programs (Dissertation)
John Hannan
MS-CIS-91-09
LINC LAB 191

In this dissertation we study a higher-order intuitionistic logic used as a specification language for a variety of tasks that treat functional programs as data objects. Such meta-programming tasks offer unique challenges including the representation of programs as data objects and the analysis of these objects. We present a technique, inspired by natural semantics and structural operational semantics, for specifying properties of programs. Specifications of this sort are presented as sets of inference rules and are encoded as clauses in a higher-order, intuitionistic meta-logic. Programs are represented by λ-terms and many features of the language such as lexical scoping are enforced through the use of λ-abstractions. Program properties are represented as propositions over these terms and are then proved by constructing proofs in our meta-logic. The meta-logic, based on natural deduction, includes inference rules for the introduction and discharge of both hypotheses and eigenvariables. We demonstrate how these rules provide simple and elegant manipulations of bound variables in functional programs. We also demonstrate how transforming proofs and proof systems in this setting provides a means for transforming meta-programs, producing new meta-programs that have certain properties or behaviors.

We argue the following points regarding these specifications and their proofs: (i)the specifications of numerous meta-programming tasks are clear, concise and well structured, providing them with simple explanations and

correctness proofs; (ii)a wide variety of meta-programming tasks can be specified in a single unified framework, and thus we can investigate and understand the relationship between various tasks; (iii)proofs describing computations or other kinds of manipulations provide a structure that can be analyzed, using established techniques from proof theory; (iv)specification in our logic have a direct translation to programs in the logic programming language λ Prolog and this translation provides a mechanism for producing experimental implementations of our meta-programs.

TRACS Users Manual and Software Reference Guide

Eric Paljug

MS-CIS-91-10
GRASP LAB 254

The Two Robotic Arm Coordination System (TRACS) of the GRASP Lab is designed to perform experiments in dynamic two arm control. The system is comprised of two PUMA 250 robot arms with modified controllers, a PA-AT host computer and an AMD 29000 high speed floating point processor board. This manual describes the system software architecture and the software interfaces between the system elements. It is intended to aid in developing software for the system.

Type-Raising and Directionality In Combinatory Grammar

Mark Steedman

MS-CIS-91-11
LINC LAB 192

The form of rules in combinatory categorial grammars (CCG) is constrained by three principles, called "adjacency", "consistency" and "inheritance". These principles have been claimed elsewhere to constrain the combinatory rules of composition and type raising in such a way as to make certain linguistic universals concerning word order under coordination follow immediately. The present paper shows that the three principles have an extremely natural expression in a unification-based interpretation of CCG in which directional information is an attribute of the arguments of functions grounded in string position. The aforementioned universals can thereby be derived as consequences of more elementary assumptions. some desirable results follow, concerning type-raising and the parser.

Surface Structure, Intonation and Meaning In Spoken Language

Mark Steedman

MS-CIS-91-12
LINC LAB 193

The paper briefly reviews a theory of intonational prosody and its relation syntax, and to certain oppositions of discourse meaning that have variously bee n called "topic and comment" "theme and rheme", "given and new", or "presupposition and focus". The theory, which is based on Combinatory Categorial Grammar, is presented in full elsewhere. the present paper examines its consequences for the automatic synthesis and analysis of speech.

A Simple, Yet Probabilistically Tractable Algorithm For First Principles Diagnosis

Ron Rymon

MS-CIS-91-13
LINC LAB 194

There are three parts to this paper. First, I present what I hope is a conclusive, worst-case, complexity analysis of two well-known formulations of the Minimal Diagnosis problem – those of [Reiter 87] and [Reggia et al 85].

I then show that Reiter's conflict-sets solution to the problem decomposes the single exponential problem into two problems, *each* exponential, that need be solved *sequentially*. From a worst case perspective, this only amounts to a factor of two, in which case I see no reason to prefer it over a simple generate-and-test approach. This is only emphasized with the results of the third part of the paper.

Here I argue for a different perspective on algorithms, that of expected, rather than worst-case performance. From that point of view, a sequence of two exponential algorithms has lesser *probability* to finish early than a single such algorithm. I show that the straightforward generate-and-test approach may in fact be somewhat attractive as it has high probability to conclude in a polynomial time, given a random problem instance.

Common Knowledge: A Survey

Marilyn A. Walker

MS-CIS-91-14
LINC LAB 195

This paper discusses the motivation behind common knowledge. Common knowledge has been argued to be necessary for joint action in general and for language use as a particular kind of joint action. However, this term has been broadly interpreted. Two major issues must be addressed: (1) What mental state corresponds to common knowledge, ie. is knowledge, belief or supposition the appropriate mental attitude? (2) What inference process allows agents to achieve common knowledge?

Most generally, common knowledge is used to describe the knowledge that is evidenced in reflexive reasoning. The term has also been used to refer to facts or objects which are

mutually salient. One of the main problems for a theory of common knowledge is whether knowledge is the appropriate mental attitude. It seems as though probabilistic beliefs might approximate the cognitive phenomenon of common knowledge more closely than knowledge.

The main problem with a usable notion of common knowledge is that inference must play a critical role in what becomes common knowledge. I discuss the nature of conversational inference. It has a number of properties that distinguish it from other inferential systems, such as being apparently abductive and probabilistic, but a precise characterization of it is an unsolved problem. I suggest that in cases where ensuring common knowledge really matters, participants in dialogue accomplish this is by exploiting opportunities for redundancy in conversation.

Structure-Based Animation Of The Human Face

Stephen M. Platt
Aaron T. Smith
Francisco Azuola
Norman I. Badler
Catherine Pelachaud

MS-CIS-91-15
GRAPHICS LAB 38

The face is an interesting object to animate for several reasons: it is an important channel of communication and therefore important to an human body animation, and it is a complex object in that it is composed of many nonrigid interacting nonarticulated regions. In this paper, we examine the face, and present it as a hierarchically structured regionally defined object. Based on this regional decomposition, and a set of primitive actions, we describe an encoding of a large set of high level facial action descriptors. We also present an application which studies that interaction between intonation and facial expressions for a given emotion. It offers a higher level of representation of the action units by grouping them into specialized functions (lips shape for phonemes, eyebrow movements). An animation system linked to facial motion property is also presented.

Dynamic Binding Communication Mechanism

Jeffrey S. Aaronson

MS-CIS-91-16
LINC LAB 196

Shastri & Ajjanagadde have proposed a biologically plausible connectionist rule-based reasoning system (hereafter referred to as a knowledge base, or KB), that represents a dynamic binding as the simultaneous, or in-phase, activity of the appropriate nodes [9]. This paper makes the first attempt at designing a biologically plausible connectionist interface mechanism between 2 distinct phase-based KB, as the next step toward providing a computational account of common-sense reasoning. The Dynamic Binding Communication Mechanism (DBCM) extracts a dynamic binding from a source KB and incorporates the binding into a destination KB so that it is consistent with the knowledge already represented in the latter. DBCM consists of several distinct, special-purpose modules. The Binding Memory (BM) is made up of several identical banks of nodes. Each time a temporally-encoded dynamic binding is extracted from the source KB, it is transferred into on of the banks, where the binding is converted to a spatially-encoded representation. The Phase Database (PD) monitors the target KB. The Phase Allocator (PA) synthesizes information from the Phase Database and from the target KB to determine the phase in which to introduce the new dynamic binding into the target KB. In turn, the PA extracts a single binding from one of the banks in the BM and introduces it into the target KB. The interface also utilizes 2 searchlight mechanisms: the first governs which bank in the BM receives binding; the second mediates between the active banks (those which are currently representing bindings), and the Phase Allocator.

The Hughes Array Co-Processor and Its Application To Robotics

Craig Sayers

MS-CIS-91-17
GRASP LAB 255

This report describes the results of twelve months research involving the Hughes array co-processor. This work began with the testing and debugging of the existing system, continued with the development of software to interface the co-processor to a host machine and concluded with the implementation of a trajectory planning algorithm for redundant manipulators.

A loader program has been developed which allows simple programs to be executed. A library of C-callable routines has also been created and this enables the fabrication of more complex systems which require a high level of interaction between the host and co-processor.

Routines to perform square root, sine and cosine functions have been designed and these have been used successfully in the development of a trajectory planning algorithm. This algorithm uses the co-processor to compute in parallel a large number of forward kinematic solutions and by doing so is able to convert a cartesian space trajectory into a joint space path for a redundant manipulator.

The performance of the processor has been analyzed and a number of recommendations have been made concerning

future implementations.

Analysis Of Dynamic Congestion Control Protocols: A Fokker-Planck Approximation

Amarnath Mukherjee

John C. Strikwerda

MS-CIS-91-18

DISTRIBUTED SYSTEMS LAB 5

We present an approximate analysis of a queue with dynamically changing input rates that are based on implicit on explicit feedback. This is motivated by recent proposals for adaptive congestion control algorithms [RaJa 88, Jac 88], where the sender's window size at the transport level is adjusted based on perceived congestion level of a bottleneck node. We develop an analysis *methodology* for a simplified system; yet it is powerful enough to answer the important questions regarding stability, convergence (or oscillations), fairness and the significant effect that delayed feedback plays on performance. Specifically, we find that, in the absence of feedback delay, the linear increase/exponential decrease algorithm of Jacobson and Ramakrishnan-Jain [Jac 88, RaJa 88] is *provably* stable and fair. Delayed feedback on the other hand, introduces oscillations for *every* individual user as well as unfairness across those competing for the same resource. While the simulation study of Zhang [Zha 89] and the fluid-approximation study of Bolot and Shanker [BoSh 90] have observed the oscillations in cumulative queue length and measurements by Jacobson [Jac 88] have revealed some of the unfairness properties, the *reasons* for these have not been identified. We identify *quantitatively* the *cause* of these effects, via-a-vis the systems parameters and properties of the algorithm used.

The model presented is fairly general and can be applied to evaluate the performance of a wide range of feedback control schemes. It is an extension of the classical Fokker-Planck equation. Therefore, it addresses traffic viability (to some extent) that fluid approximation techniques do not address.

Programming With *Jack*

Cary B. Phillips

MS-CIS-91-19

GRAPHICS LAB 39

This manual describes the implementation of *Jack* with emphasis on how to extend it and modify it. the principle purpose of this manual is to describe what functions in the *Jack* libraries are available to be used in writing new features for *Jack*. the manual also gives an overview of how *Jack* works, for those interested in modifying its current behavior. This manual assumes that you already know how to use *Jack*, and are familiar with its basic terminology.

GRASP NEWS Volume 7, Number 1 Spring 1991

Various Contributors

MS-CIS-91-20

GRASP LAB 256

Since its beginning in 1983, the *GRASP News* has chronicled the research efforts of the Grasp Laboratory. This edition, which covers developments for the year 1990, follows the format of previous editions. The Feature Article, however, in a departure from tradition does not highlight a particular research project. The research abstract summarize the progress of students, postdoctoral fellows, visiting researchers, and faculty. The abstracts are classified into three different areas – *Visions Research, Robotics Research,* and *Distributed Real-Time Systems Research.* there is also a section on laboratory *Software and Hardware Developments.* this edition comprises 40 articles from 45 contributors, which makes it the largest *GRASP News* ever!

Representing Powerdomain Elements As Monadic Second Order Predicates

Carl A. Gunter

MS-CIS-91-21

LOGIC & COMPUTATION 28

This report characterizes the powerdomain constructions which have been used in the semantics of programming languages in terms of formulas of first order logic under a preordering of provable implication. This provides an intuitive representation which suggests a new form of powerdomain—called the *mixed* powerdomain—which expresses data in a different way from the well-known constructions from programming semantics. It can be shown that the mixed powerdomain has many of the properties associated with the convex powerdomain such as the possibility of solving recursive equations and a simple algebraic characterization.

Tree-Adjoining Grammars and Lexicalized Grammars

Aravind K. Joshi

Yves Schabes

MS-CIS-91-22

LINC LAB 197

In this paper, we will describe a tree generating system called tree-adjoining grammar (TAG) and state some of the recent result about TAGs. The work on TAGs is motivated by linguistic considerations. However, a number of formal results have been established for TAGs, which we believe, would be of interest to researchers in tree grammars and tree automata. After giving a short introduction to TAG, we

briefly state these result concerning both the properties of the string sets and tree sets. We will also describe the notion of lexicalization of grammars and investigate the relationship of lexicalization to context-free grammars (CFGs) and TAGs.

An Abstract Interpretation For ML Equality Kinds

Carl A. Gunter
(University of Pennsylvania)
Elsa L. Gunter
(AT&T Bell Laboratories)
David B. MacQueen
(AT&T Bell Laboratories)

MS-CIS-91-23
LOGIC & COMPUTATION 29

The definition of Standard ML provides a form of generic equality which is inferred for certain types, called *equality types*, on which it is possible to define a computable equality relation. However, the standard definition is incomplete in the sense that there are interesting and useful types which are *not* inferred to be equality types but which nevertheless have a computable equality relation. In this paper we introduce a refinement of the Standard ML system of equality types and prove that our system is sound and *complete* with respect to the existence of a computable equality. Our technique is based on an abstract interpretation of ML operators as monotone functions over a three point lattice. We show how the equality relation can be defined (as an ML program) from the definition of a type with our equality property. We then demonstrate a sound, efficient algorithm for inferring the equality property which corrects the limitations of the standard definition in all cases of practical interest.

Unification Of Simply Lambda-Terms As Logic Programming

Dale Miller

MS-CIS-91-24
LINC LAB 198

The unification of simply typed lambda-terms modulo the rules of beta- and eta-conversions is often called "higher-order" unification because of the possible presence of variables of functional type. This kind of unification is undecidable in general and if unifiers exist, most general unifiers may not exist. In this paper, we show that such unification problems can be coded as a query of the logic programming language L-lambda in a natural and clear fashion. In a sense, the translation only involves explicitly axiomatizing in L-lambda the notions of equality and substitution of the simply typed lambda-calculus: the rest of the unification process can be viewed as simply an interpreter of L-lambda searching for proofs using those axioms.

[[Appears in the Proceedings of the 1991 International Conference on Logic Programming, edited by Koichi Furukawa, June 1991.]]

Unification-Based Tree Adjoining Grammars

K. Vijay-Shanker
(University of Delaware)
Aravind K. Joshi
(University of Pennsylvania)

MS-CIS-91-25
LINC LAB 199

Many current grammar formalisms used in computational linguistics take a unification-based approach that use structures (called feature structures) containing sets of feature-value pairs. In this paper, we describe a unification-based approach to Tree Adjoining Grammars (TAG). The resulting formalism (UTAG) retains the principle of factoring dependencies and recursion that is fundamental to TAGs (see[Schabes*et. al.*, 1988]). We give some linguistic examples using UTAG and informally discuss the descriptive capacity of UTAG, comparing it with other unification-based formalisms. Finally, based on the linguistic theory underlying TAGs, we propose some stipulations that can be placed on UTAG grammars. In particular, we stipulate that the feature structures associated with the nodes in an elementary tree are bounded (there is an analogous stipulation on GPSG). Grammars that satisfy these stipulations are equivalent to TAG. Thus, even with these stipulations, UTAGs have more power than CFG-based unification grammars with the same stipulations.

[[To Appear in *Unification-Based Grammars" (ed. Jurgen Wedekind), MIT PRESS, 1991*]]

Relevant Consequence and Empirical Inquiry

Daniel N. Osherson
(IDIAP)
Scott Weinstein
(University of Pennsylvania)

MS-CIS-91-26
LOGIC & COMPUTATION 30

A criterion on adequacy is proposed for theories of relevant consequence. According to the criterion, scientists whose deductive reasoning is limited to some proposed subset of the standard consequence relation must not thereby suffer a reduction in scientific competence. A simple theory

of relevant consequence is introduced and show to satisfy the criterion with respect to a formally defined paradigm of empirical inquiry.

Occlusions As A Guide For Planning The Next View

Jasna Maver
Ruzena Bajcsy

MS-CIS-91-27
GRASP LAB 257

The task of constructing a volumetric description of a scene from a single image is an underdetermined problem, whether it is a range or an intensity image. To resolve the ambiguities that are caused by occlusions in images, we need to take sensor measurements from several different views. We have limited ourselves to range images obtained by a laser scanning system. It is an active system which can encounter two types of occlusions. An occlusion arises either when the reflected laser light does not reach the camera or when the direct laser light does not reach the scene surface. The task of 3-D data acquisition is divided into two subproblems: to acquire the depth information from one scanning plane and to select the proper scanning planes from which the direct laser light illuminates the entire scene. The first kind of occlusions (range shadows) are easily detected and can be used in designing an efficient algorithm. We develop a strategy to determine the sequence of different views using the information in a narrow zone around the occluded regions. Occluded regions are approximated by polygons. Based on the height information of the border of the occluded regions and geometry of the edges of the polygonal approximation, the next views in the same scanning plane are determined. From the acquired information in the first scanning plane the directions of the next scanning planes for further data acquisition are computed.

Verb Composition For The Animation Of Natural Language Instructions

Libby Levinson

MS-CIS-91-28
GRAPHICS LAB 40
LINC LAB 200

This report is an investigation of issues encountered in generating a short simulation from a set of instructions. A method for specifying simulations at a task-level, rather than by individual motion, is discussed. The research was conducted using a set of instructions that describe the removal of a Fuel Control Valve from an aircraft.

Logical And Computational Aspects Of Programming With Sets/Bags/Lists

Val Breazu-Tannen
Ramesh Subrahmanyam

MS-CIS-91-29
LOGIC & COMPUTATION 31

We study issues that arise in programming with primitive recursion over non-free datatypes such as lists, bags and sets. Programs written in this style can lack a meaning in the sense that their outputs may be sensitive to the choice of input expression. We are, thus, naturally lead to a set-theoretic denotational semantics with partial functions. We set up a logic for reasoning about the definedness of terms and a deterministic and terminating evaluator. The logic is shown to be sound in the model, and its recursion free fragment is shown to be complete for proving definedness of recursion free programs. The logic is then shown to be as strong as the evaluator, and this implies that the evaluator is compatible with the provable equivalence between different set (or bag, or list) expressions . Oftentimes, the same non-free datatype may have different presentations, and it is not clear *a priori* whether programming and reasoning with the two presentations are equivalent. We formulate these questions, precisely, in the context of alternative presentations of the list, bag, and set datatypes and study some aspects of these questions. In particular, we establish back-and-forth translations between the two presentations, from which it follows that they are equally expressive, and prove results relating proofs of program properties, in the two presentations.

Design Of A Tool-Surrounding Compliant Instrumented Wrist

Thomas Lindsay
Richard P. Paul

MS-CIS-91-30
GRASP LAB 258

Interaction between robot and environment is an extremely important aspect of robotic research. Compliance helps reduce the effects of impact when there is robot/environment interaction. To accomplish useful tasks, it is important to implement hybrid control; accurate position control is needed in unconstrained directions and accurate force control is needed in constrained direction. Force control can be more responsive with a compliant force/torque sensor [3], but positional accuracy is reduced with compliance. An instrumented compliant wrist device can be used to achieve both responsive force control and accurate position control.

The wrist is connected in series between the end of the

robot and the tool. The wrist device uses rubber elements for compliance and damping, and a serial linkage, with potentiometers at each joint, is used for sensing the defections produced in the wrist. Several major improvements are proposed for the Xu wrist. The wrist can be designed to surround the tool, thus reducing the distance between the end of the robot and the end of the tool, thus reducing the distance between the end of the robot and the end of the tool. The compliant structure is redesigned for more even compliance, and the sensing structure kinematics are simplified.

In this over, the compliance, kinematics, and accuracy of the wrist will be presented. Also, software for finding the wrist transform, and plans for the wrist are given.

Communicating Shared Resources: A Paradigm For Integrating Real-Time Specification And Implementation

Insup Lee
Susan Davidson
Richard Gerber
MS-CIS-91-31
GRASP LAB 259

The timed behavior of distributed real-time systems can be specified using a formalism called Communicating Shared Resources, or CSR. The underlying computation model of CSR is resource-based in which multiple resources execute synchronously, while processes assigned to the same resource are interleaved according to their priorities. CSR bridges the gap between an abstract computation model and implementation environments, but is too complex to be treated as a process algebra. We therefore give a calculus for CSR (CCSR), that provides the ability to perform equivalence proofs by syntactic manipulation. We illustrate how a CSR specification can be translated into the CCSR formalism using a periodic timed producer-consumer example, and how a translated CSR specification can be shown correct using syntactic manipulations.

Logic Programming In A Fragment Of Intuitionistic Linear Logic: Extended Abstract

Joshua Hodas
Dale Miller
MS-CIS-91-32
LOGIC & COMPUTATION 32

Logic programming languages based on fragments of intuitionistic logic have recently been developed and studied by several researchers. In such languages, implications are permitted in goals and in the bodies of clauses. Attempting to prove a goal of the form $D \supset G$ in a context Γ leads to an attempt to prove the goal G in the extended context $\Gamma \cup \{D\}$. While an intuitionistic notion of context has many uses, it has turned out to be either too powerful or too limiting in several settings.

We refine the intuitionistic notion of context by using a fragment of Girard's linear logic that includes additive and multiplicative conjunction, linear implication, universal quantification, the "of course" exponential, and the constants 1 (the empty context) and $\top$ (for "erasing" contexts). After presenting our fragment of linear logic, which contains the hereditary Harrop formulas, we show that the logic has a goal-directed interpretation. We also show that the nondeterminism that results from the need to split contexts in order to prove a multiplicative conjunction can be handled by viewing proof search as a process that takes a context, consumes part of it, and returns the rest (to be consumed elsewhere). The complete specification of an interpreter for this logic is presented.

Examples taken from theorem proving, natural language parsing, and data base programming are presented: each example requires a linear, rather than intuitionistic, notion of context to be modeled adequately.

Combining A Type Hierarchy With A Rule-Based Reasoner

Lokendra Shastri
D.R. Mani
MS-CIS-91-33 LINC LAB 201

This report describes an efficient connectionist knowledge representation and reasoning system that combines rule-based reasoning with inheritance and classification within an *IS-A* hierarchy. In addition to a type hierarchy, the proposed system can encode generic facts such as 'Cats prey on birds' and rules such as 'If x preys on y then y is scared of x' and use them to infer that Tweety (who is a Canary) is scared of Sylvester (who is a Cat). The system can also encode qualified rules such as if an *animate* agent walks into a *solid object* then the agent gets hurt'. The proposed system can answer queries in time that is only *proportional* to the *length* of the shortest derivation of the query and is independent of the size of the knowledge base. The system maintains and propagates variable bindings using temporally synchronous — i.e., in-phase — firing of appropriate nodes.

Formal Models For Concurrent Communicating Systems

Anthony S. Kosky
MS-CIS-91-34
LOGIC & COMPUTATION 33

This report was originally written to fulfill in part the requirements of the author's WPE examinations, part of the qualifying examinations for the University of Pennsylvania'a Computer Science Ph.D program. The report first introduces CCS and uses it to illustrate various features of established methods of modelling concurrent, communicating systems. The report then goes on to describe and investigate two new models for such systems: The *Chemical Abstract Machine,* a simple yet predominant in most models for such systems; and the π-calculus, a calculus similar in many respects to CCS, but able to model mobile processes and other, more difficult phenomena.

RTC: Language Support For Real-Time Concurrency

Victor Wolfe
Susan Davidson
Insup Lee
MS-CIS-91-35
GRASP LAB 260

This paper presents language constructs for the expression of timing and concurrency requirements in distributed real-time programs. Our programming paradigm combines an object-based paradigm for the specification of shared resources, and a distributed transaction-based paradigm for the specification of application processes. Resources provide abstract views of shared system entities, such as devices and data structures. Each resource has a state and defines a set of *actions* that can be invoked by processes to examine or change its state. A resource also specifies scheduling constraints on the execution of its actions to ensure the maintenance of its state's consistency. Processes access resources by invoking actions and express precedence, consistency. Processes access resources by invoking actions and express precedence, consistency and timing constraints on action invocations. The implementation of our language constructs with real-time scheduling and locking for concurrency control is also described.

A Framework For Visual Observation (Dissertation Proposal)

Tarek M. Sobh
MS-CIS-91-36
GRASP LAB 261

We address the problem of observing a moving agent. In particular, we propose a system for observing a minipulation process, where a robot hand manipulates an object. A discrete event dynamic systems (DEDS) frame work is developed for the hand/object interaction over time and a stabilizing overserver is constructed. Low-level modules are developed for recognizing the "events" that causes state transitions within the dynamic manipulation system. The work examines closely the possibilities for errors, mistakes and uncertainties in the manipulation system, observer construction process and event identification mechanisms. The system utilizes different tracking techniques in order to observe and recognize the task in an *active, adaptive* and *goal-directed* manner.

The Role Of Vergence Micromovements On Depth Perception

Antônio Francisco Júnior
MS-CIS-91-37
GRASP LAB 262

A new approach in stereo vision is proposed which recovers 3D depth information using *continuous vergence angle control* with simultaneous local correspondence response. This technique relates elements with the *same* relative position in the left and right images for a continuous sequence of vergence angles. the approach considers the extremely fine vergence movements about a given fixation point within the depth of field boundaries. It allows the recovery of 3D depth information given the knowledge of the system's geometry and a sequence of pairs $[\alpha_I, \mathbf{C}_i]$, where α_i is the i^{th} vergence angle and $\mathbf{C}_i$ is the i^{th} matrix of correspondence responses. The approach has several advantages over the current ones. First, due to its local operation characteristics, the resulting algorithms can be implemented in a modular hardware scheme. Second, unlike currently use algorithms, there is no need to compute depth from disparity values; at the cost of the acquisition of a sequence of images during the micromovements. the approach also greatly reduces the errors in stereo due to the sensor quantization. Last, and most important of all, the approach is supported by experimental results from physiology and psychophysics. Physiological results show that the human eye performs fine movements during the process of fixation on a single point, which are collectively called *physiological nystagmus.* One such movement, called *binocular flicks,* happens in opposing directions and produces convergence/divergence of the eyes. These are the micromovements that we suppose are the basis for depth perception. Therefore, the approach proposes a functional correlation between these vergence micromovements, depth perception, stereo acuity and stereo fusion.

Performance Evaluation via Perturbation Analysis

Tarek M. Sobh
MS-CIS-91-38
GRASP LAB 263

In this paper we present an overview for the development of a theory for analyzing and predicting the behavior if discrete event dynamic systems (DEDS). DEDS are dynamic systems in which state transitions are caused by internal, discrete events in the system. DEDS are attracting considerable interest, current applications are found in manufacturing systems, communications and air traffic systems, future applications will include robotics, computer vision and artificial intelligence. We will discuss the perturbation analysis technique (PA) for evaluation the performance of DEDS.

Discrete Event Dynamic Systems: An Overview
Tarek M. Sobh
MS-CIS-91-39
GRASP LAB 264

In this report we present an overview for the development of a theory for discrete event dynamic systems (DEDS). Dynamic systems are usually modeled by finite state automata with partially overservable events together with a mechanism for enabling and disabling a subset of state transitions. DEDS are attracting considerable interests, current applications are found in manufacturing systems, communications and air traffic systems, future applications will include robotics, computer vision and AI. We will discuss notions of modeling, stability issues, observability, feedback and invertibility. We will also discuss the perturbation analysis technique (PA) for analyzing and describing the behavior of DEDs.

Teleprogramming: Towards Delay-Invariant Remote Manipulation (Dissertation)
Janez Funda
MS-CIS-91-40
GRASP LAB 265

This dissertation addresses the problem of remote manipulation in the presence of communication delays. Delays occur with earth-based control of a robotic system in space or when an untethered submersible system is controlled from the surface via an acoustic communication channel. The resulting delay in obtaining position and force feedback from the remote slave arm(s) makes direct teleoperation infeasible.

We propose a new control methodology, called *teleprogramming*, which allows for efficient control for a robotic system in the presence of significant feedback delays without substantial degradation in the overall system performance. A *teleprogramming* system allows the operator to kinesthetically, as well as visually, interact with a graphical simulation of the remote environment and to interactively, online teleprogram the remote manipulator through a sequence of elementary symbolic instructions. These instructions are generated automatically by the operator's station software in real time as the task progresses. The slave robot executes these symbolic commands delayed in time and, should an error occur, allows the operator to specify the necessary corrective actions and continue with the task.

Teleprogramming offers a practical compromise between the ultimate and the feasible, and provides an effective and time-efficient approach to remote manipulation. Advantages of *teleprogramming* over existing control methodologies include a relatively modest required level of remote site autonomy, and the absence of the need of complex automatic task planners and preprogrammed error recovery modules.

This document describes the overall conceptual architecture of *teleprogramming* and presents a detailed treatment of all major components of *teleprogramming* system. An operational prototype system is described an preliminary experimental results are reported. Experimental results have confirmed the validity and feasibility of the *teleprogramming* control methodology. Sustained and efficient remote control of a robot manipulator in the presence of a five second feedback delay was successfully accomplished for simple contact tasks.

A Comparison Of Compressed and Uncompressed Transmission Modes
Tarek M. Sobh
Jaffar Rehman
MS-CIS-91-41
GRASP LAB 266

In this paper we address the problem of host to host communication. In particular, we discuss the issue of efficient and adaptive transmission mechanisms over possible physical links. We develop a tool for making decisions regarding the flow of control sequences and data from and to a host. This issue of compression is discussed in details, a decision box and an optimizing tool for finding the appropriate thresholds for a decision are developed. Physical parameters like the data rate, bandwidth of the communication medium, distance between the hosts, abaud rate, levels of discretization, signal to noise ration and propagation speed of the signal are taken into consideration while developing our decision system. Theoretical analysis is performed to develop mathematical models for the optimization algorithm. Simulation models are also developed for testing both the optimization and the decision tool box.

Generation and Synchronous Tree-Adjoining Grammars

Stuart M. Sheiber
(Harvard University)
Yves Schabes
(University of Pennsylvania)

MS-CIS-91-42
LINC LAB 202

Synchronous Tree-Adjoining Grammars

Stuart M. Sheiber
(Harvard University)
Yves Schabes
(University of Pennsylvania)

MS-CIS-91-43
LINC LAB 203

The unique properties of tree-adjoining grammars (TAG) present a challenge for the application of TAGs beyond the limited confines of syntax, for instance, to the task of semantic interpretation or automatic translation of natural language. We present a variant of TAGs, called synchronous TAGs, which characterize correspondences between languages. The formalism's intended usage is to relate expressions of natural languages to their associated semantics represented in a logical form language, or to their translates in another natural language; in summary, we intend it to allow TAGs to be used beyond their role in syntax proper. We discuss the application of synchronous TAGs to concrete examples, mentioning primarily in passing some computational issues that arise in its interpretation.

Using Lexicalized Tags For Machine Translation

Anne Abeillé
(University of Paris)
Yves Schabes
(University of Pennsylvania)
Aravind K. Joshi
(University of Pennsylvania)

MS-CIS-91-44
LINC LAB 204

Lexicalized Tree Adjoining Grammar (LTAG) is an attractive formalism for linguistic description mainly because of its extended domain of locality and it s factoring recursion out from the domain of local dependencies (Joshi, 19845, Kroch and Joshi, 1985, Abeillé, 1988). LTAG's extended domain of locality enables on to localize syntactic dependencies (such as filler-gap), as well as semantic dependencies (such as predicate-arguments). The aim of this paper is to show that these properties combined with the lexicalized property of LTAG are especially attractive for machine translation.

The transfer between two languages, such as French and English, can be done by putting directly into correspondence large elementary universe without going through some interlingual representation and without major changes to the source and target grammars. The underlying formalism fro the transfer is "synchronous Tree Adjoining Grammars" (Sheiber and Schabes [1990]). Transfer rules are stated as correspondences between nodes of trees of large domain of locality which are associated with words. We can thus define lexical transfer rules that avoid the defects of a mere word-to-word approach but still benefit from the simplicity and elegance of a lexical approach.

We rely on the French and English LTAG grammars (Abeillé [1988], Abeillé [1990(b)], Abeillé et al. [1990], Abeillé and Schabes [1989, 1990]) that have been designed over the past two years jointly at University of Pennsylvania and University of Paris 7-Jussieu.

Surface and Volumetric Segmentation Of Complex 3-D Objects Using Parametric Shape Models (Dissertation)

Alok Gupta

MS-CIS-91-45
GRASP LAB 267

The problem of part definition, description, and decomposition is central to the shape recognition systems. In this dissertation, we develop an integrated framework for segmenting dense range data of complex 3-D scenes into their constituent parts in terms of surface and volumetric primitives. Unlike previous approaches, we use geometric properties derived from surface, as well as volumetric models, to *recover* structured descriptions of complex objects without *a priori* domain knowledge or stored models.

To recover shape descriptions, we use *bi-quadric* models for surface representation and *superquadric* models for object-centered volumetric representation. The surface segmentation uses a novel approach of *searching* for the best piecewise description of the image in terms of bi-quadric ($z = f(x, y)$) models. It is used to generate the region adjacency graphs, to localize surface discontinuities, and to derive global shape properties of the surfaces. A superquadric model is recovered for the entire data set and residuals are computed to evaluate the fit. The goodness-of-fit value based on the inside-outside function, and the mean-squared distance of data from the model provide quantitative evaluation of the model. The qualitative evaluation criteria check the local consistency of the model in the form of residual

maps of *overestimated* and *underestimated* data regions.

The control structure invokes the models in a systematic manner, evaluates the intermediate descriptions, and integrates them to achieve final segmentation. Superquadric and bi-quadric models are recovered in parallel to incorporate the best of the coarse-to-fine and fine-to-coarse segmentation strategies. The model evaluation criteria determine the dimensionality of the scene, and decide whether to terminate the procedure, or selectively refine the segmentation by following a global-to-local part segmentation approach. The control module generates hypotheses about superquadric models at clusters of underestimated data and performs controlled extrapolation of the part-model by shrinking the global model. As the global model shrinks and the local models grow, they are evaluated and tested for termination or further segmentation.

We present results on real range images of scenes of varying complexity, including objects with occluding parts, and scenes where surface segmentation is not sufficient to guide the volumetric segmentation. We analyze the issue of segmentation of complex scenes thoroughly by studying the effect of missing data on volumetric model recovery, generating object-centered descriptions, and presenting a complete set of criteria for the evaluation of the superquadric models. We conclude by discussing the applications of our approach in data reduction, 3-D object recognition, geometric modeling, automatic model generation, object manipulation, and active vision.

Self Organizing Feature Maps and Their Applications To Robotics

Craig Sayers

MS-CIS-91-46
GRASP LAB 268

The self-organizing feature maps developed by Kohonen appear to capture some of the advantages of the natural systems on which they are based. A summary of the operation of this form of artificial neural network is presented. It was concluded that the primary benefits of using self-organizing feature maps result from their adaptability and plasticity while most problems are largely caused by the lack of a rigorous mathematical foundation.

Two different robotics applications are described. In the first, developed by Martinez and Schulten, a hierarchical structure composed of many self-organizing feature maps is used to control a five degree of freedom robot arm. While it was noted that there may be some practical problems, the general idea of using a hierarchical structure appears sound and may be applicable to a wider range of problems.

The second robotics application was developed by Saxon and Mukherjee. They used a single self-organizing feature map to learn the motion map of a two degree of freedom arm. The use of such a system should simplify path planning by combining multiple constraints into a 2-D structure.

Optimal Randomized Algorithms For Multipacket and Wormhole Routing On the Mesh

Sanguthevar Rajasekaran
Mukund Raghavachari

MS-CIS-91-47
GRASP LAB 269

In this paper, we present a randomized algorithm for the multipacket (i.e., $k - - - k$) routing problem on an nxn mesh. the algorithm competes with high probability in at most $kn + O(k \log n)$ parallel communication steps, with a constant queue size of $O(k)$. the previous best known algorithm [4] takes $\frac{5}{4}kn + O(\frac{kn}{f(n)})$ steps with a queue size of $O(kf(n))$ (for any $1 \leq f(n) \leq n$). We will also present a randomized algorithm for the wormhold model permutation routing problem for the mesh that completes in at the most $kn + O(k \log n)$ steps, with a constant queue size of $O(k)$, where k is the number of flits that each packet is divided into. The previous best result [6] was also randomized and had a time bound of $kn + O(\frac{kn}{f(n)})$ with a queue size of $O(kf(n))$ for any $1 \leq f(n) \leq n$). the two algorithms that we will present are optimal with respect to queue size. The time bounds are within a factor of two of the only known lower bound.

Analysis and Simulation Of Mechanical Systems With Multiple Frictional Contacts

Yin Tien Wang
Vijay R. Kumar

MS-CIS-91-48
GRASP LAB 270

In many engineering applications such as assembly of mechanical components, robot manipulation, gripping, fixturing and part feeding, there are situations in which a rigid body is subject to multiple frictional contacts with other bodies. It is proposed to develop a systematic method for the analysis and simulation of such systems. A detailed study is presented on rigid body impact laws, and the assumption of contact compliance is investigated.

Theoretical Aspects Of Schema Merging

Peter Buneman
Susan Davidson
Anthony Kosky

MS-CIS-91-49
LOGIC & COMPUTATION 34

A general technique for merging database schemas is developed that has a number of advantages over existing techniques, the most important of which is that schemas are placed in a partial order that has bounded joins. This means that the merging operation, when it succeeds, is both associative and commutative, i.e., that the merge of schemas is independent of the order in which they are considered – a property not possessed by existing methods. The technique is interactive in that users made assertions about the relationships between the nodes of the schemas to be merged. These assertions are then considered to be elementary schemas, and are combined with the schemas using precisely the same merging operation.

The technique is general and can be applied to a variety of data models. It can also deal with certain cardinality constraints that arise through the imposition of keys. A prototype implementation, together with a graphical interface, has been developed.

Natural Language Control Of Animation Of Task Performance In A Physical Domain (Dissertation)

Jugal Kumar Kalita

MS-CIS-91-50
GRAPHICS LAB 41

We establish a link from natural language statements describing actions to be performed by an agent to a semantic representation suitable for achieving effective control of a computer-driven graphical animation system. A representation scheme based on decompositional analysis is developed emphasizing the requirement that algorithmic implementability of the underlying semantic primitives is our primary concern. Our primitives pertain to mechanical characteristics of the "kernel" tasks denoted by a class of action verbs (verbs whose underlying tasks deal with an agent manipulating one or more objects); they refer to geometric constraints and goals that need to be achieved, kinematic and dynamic characteristics, and certain aspectual characteristics such as repetitiveness of one or more sub-actions, definedness of termination points, etc. We provide lexical entries for a few verbs in terms of such primitives.

We also analyze the manner in which prepositional and adverbial modifiers affect the representation as well as the execution of the basic actions denoted by the verbs. Such modifiers either provide values of arguments for the verbs' internal representations, modify default argument values, or provide values of non-obligatory arguments. We obtain semantic representations for a few prepositions and adverbs in a fashion integrable into the scheme for verbal meaning representation. We have developed a system to demonstrate the validity of the results obtained; such a system establishes channels of communication with existing animation software developed at the Graphics Laboratory at the University of Pennsylvania.

On Call Migration

Ming Chit Tam

MS-CIS-91-51
DISTRIBUTED SYSTEMS LAB 6

In an environment where network resources are reserved e.g, telephone networks, the path with smallest number of hops is pre- fered and other alternate paths are used only when there the shortest path is full. However if the alternate path is longer more network resources are devoted to the circuit and this in turn could worsen the situation. Circuit migration is a solution to reduce the amount of resources inefficiently used due to al- ternate routing in connection oriented networks. By rerouting a circuit when its shortest path becomes available, one can smooth out the congestion and increases the utilization of the network.

The overhead of circuit migration is comparable to call set up and the tradeoff of circuit migration is improvement in perfor- mance vs. some additional call processing capacity. In this re- port we will focus on the above tradeoff, evaluating it analyti- cally and by simulation on a completely connected topology.

Our initial results indicate that migration could improve the performance of the network at high load but it has to be done very often. Such a large amount of overhead could be expensive enough to offset the gain in performance. On further investiga- tion, we discover that threshing can also occur in circuit migra- tion. We proposed two solutions to the problem. The first solu- tion is to migrate only when the shortest path is no longer high- ly utilized. The second solution migrates a circuit only if its path is congested. A hybrid solution using the two above is also examined. We will also address the reordering problem that could occur when a circuit is transferred to a new path.

Fast Algorithms For Generating Discrete Random Variates With Changing Distribution

Sanguthevar Rajasekaran
Keith W. Ross

MS-CIS-91-52
GRASP LAB 271

One of the most fundamental and frequently used operations in the process of simulating a stochastic discrete event system is the generation of a nonuniform discrete random

variate. The simplest form of this operation can be stated as follows: Generate a random variable X which is distributed over the integers $1, 2, \ldots n$ such that $P(X = i) = p_i$. A more difficult problem is to generate X when the p_i's change with time. For this case, there is a well-known algorithm which takes $O(\log n)$ time to generate each variate. Recently Fox [4] presented an algorithm that takes an expected $o(\log n)$ time to generate each variate under assumptions restricting the way the p_i's can change.

In this paper we present algorithm for discrete random variate generation that take an expected $O(1)$ time to generate each variate. Furthermore, our assumptions on how the p_i's change are less restrictive than those of Fox. The algorithms are quite simple and can be fine-tuned to suit a wide variety of application. The application to the simulation of queueing networks is discussed in some detail.

Dynamic Time Windows and Generalized Closed-Loop/Open-Loop Mechanisms For Congestion Control Of Data Traffic In High Speed Wide Area Networks

Amarnath Mukherjee
(University of Pennsylvania)
Lawrence H. Landweber
(University of Wisconsin)
Theodore Faber
(University of Wisconsin)

MS-CIS-91-53
DISTRIBUTED SYSTEMS LAB 7

This paper presents a set of mechanisms for congestion control of data traffic in high speed wide area networks (HSWANs) along with preliminary performance results. The model of the network assumes reservation of resources based on average requirements. The mechanisms address (a) the different network time constants (short term and medium-term), (b) admission control that allows controlled variance of traffic as a function of medium-term congestion, and (c) prioritized scheduling which is based on a new fairness criterion. This latter criterion is perceived as the appropriate fairness measure for HSWANs.

Preliminary performance studies show that the queue length statistics at switching nodes (mean, variance and max) are approximately proportional to the end-point 'time window' size. Further,

- when network utilization approaches unity, the time window mechanism can protect the network from buffer overruns and excessive queueing delays, and
- when network utilization level is smaller, the time window may be increased to allow a controlled amount of variance that attempts to simultaneously meet the performance goals of the end-user and that of the network.

The prioritized scheduling algorithms proposed and studied in this paper are a generalization of the Virtual Clock algorithm [Zhang 1989]. The study here investigates

- necessary and sufficient conditions for accomplishing desired fairness,
- simulation and (limited analytical results for expected waiting times,
- ability to protect against misbehaving users, and
- relationship between end-point admission control (Time-Window) and internal scheduling ('Pulse' and Virtual Clock) at the switch

Flexible Support For Trauma Management Through Goal-Directed Reasoning and Planning

Bonnie L. Webber
Ron Rymon
John R. Clarke

MS-CIS-91-54
LINC LAB 205

REMINDER: ****PAPER TO FOLLOW BONNIE. WHEN RECEIVED, DELETE THIS MESSAGE****

We describe a system, *TraumAID*, which has been designed to provide decision support throughout the initial definitive management of severely injured patients (i.e., after their initial evaluation, resuscitation, and stabilization). Over the course of initial definitive management, TraumAID recommends appropriate procedures to be carried out, based on currently available evidence and on the complexity and urgency of the situation. TraumAID's ability to deal flexibly with complex and often urgent situations comes from its ability to reason separately about the management *goals* that should be achieved and about the *means* that are situationally appropriate for achieving them.

In this paper, we describe TraumAID's approach to trauma management in more detail, showing in particular how it enables TraumAID to adapt its reasoning and recommendations to the *urgency* with which a patient's condition must be addressed.

Fast Algorithms For Generating Discrete Random Variates With Changing Distribution

Sanguthevar Rajasekaran
Keith W. Ross
MS-CIS-91-52
GRASP LAB 271

One of the most fundamental and frequently used operations in the process of simulating a stochastic discrete event system is the generation of a nonuniform discrete random variate. The simplest form of this operation can be stated as follows: Generate a random variable X which is distributed over the integers $1, 2, \ldots n$ such that $P(X = i) = p_i$. A more difficult problem is to generate X when the p_i's change with time. For this case, there is a well-known algorithm which takes $O(\log n)$ time to generate each variate. Recently Fox [4] presented an algorithm that takes an expected $o(\log n)$ time to generate each variate under assumptions restricting the way the p_i's can change.

In this paper we present algorithm for discrete random variate generation that take an expected $O(1)$ time to generate each variate. Furthermore, our assumptions on how the p_i's change are less restrictive than those of Fox. The algorithms are quite simple and can be fine-tuned to suit a wide variety of application. The application to the simulation of queueing networks is discussed in some detail.

Flexible Support For Trauma Management Through Goal-Directed Reasoning and Planning

Bonnie L. Webber
Ron Rymon
John R. Clarke
MS-CIS-91-54
LINC LAB 205

We describe a system, *TraumAID*, which has been designed to provide decision support throughout the initial definitive management of severely injured patients (i.e., after their initial evaluation, resuscitation, and stabilization). Over the course of initial definitive management, TraumAID recommends appropriate procedures to be carried out, based on currently available evidence and on the complexity and urgency of the situation. TraumAID's ability to deal flexibly with complex and often urgent situations comes from its ability to reason separately about the management *goals* that should be achieved and about the *means* that are situationally appropriate for achieving them.

In this paper, we describe TraumAID's approach to trauma management in more detail, showing in particular how it enables TraumAID to adapt its reasoning and recommendations to the *urgency* with which a patient's condition must be addressed.

Supporting Real-Time Concurrency

Victor Wolfe
MS-CIS-91-55
GRASP LAB 272
DISTRIBUTED SYSTEMS LAB 8

Concurrent real-time applications are complicated since both timing and consistency constraints must be met for correct performance. Furthermore, techniques to enforce these two forms of constraints are often incompatible. For instance, priority-driven preemptive scheduling, which is optimal for meeting timing constraints in some systems, may leave a shared resource's state inconsistent. On the other hand, mutual exclusion techniques that ensure the consistency of shared resources are not well-suited to meeting timing constraints. This dissertation develops concepts and programming language constructs for facilitating the enforcement of both real-time and consistency constraints in applications with concurrency.

Our programming paradigm combines an object-based paradigm for the specification of shared resources, and a distributed transaction-based paradigm for the specification of application processes. Resources provide abstract views of shared system entities, such as devices and data structures. Each resource has a state and defines a set of *actions* that can be invoked by processes to examine or change its state. A resource also specifies scheduling constraints on the execution of its actions to ensure the maintenance of its state's consistency. Processes access resources by invoking actions and express precedence, consistency and timing constraints on action invocations. The implementation of our language constructs with real-time scheduling and locking for concurrency control is also described, including a novel deadlock prevention technique. The utility of the constructs are demonstrated in two ways. First, we describe their use to solve a general concurrent real-time problem called *timed atomic commitment*. We then describe how they were used to program a graphic simulation of two robot arms coordinating to pick up a moving object under timing constraints.

Three Dimensional Workspace Visualization For Redundant Articulated Chains

Tarek Alameldin
MS-CIS-91-56
GRAPHICS LAB 42

This thesis deals with the problem of the 3D workspace visualization for anthropomorphic linkages, not only those with redundant degrees of freedom but also those with joint

limits. Although the workspace problem has important applications in a variety of fields such as computer-aided design, ergonomic studies, and robotics, the problem's computational complexity has never been analyzed. In addition, previous techniques suffer from one or more of the following drawbacks: high computational cost, computing 2D workspace cross sections, dealing with manipulators that have specialized geometry, or sensitivity to geometrical and numerical errors or approximations.

We analyze the computational complexity of the problem and prove that it is NP-hard. Then, we decompose it into three major subproblems: workspace point generation, visualization, and criteria selection. We describe and compare different techniques for computing workspace points: direct kinematics based algorithms, nonlinear programming based algorithms, and force application based algorithms. Each class of these algorithms has advantages and disadvantages, and none of them supersedes the others in all applications. Instead of debating the merits of these algorithms, we integrate them into "Hybrid algorithms" that are capable of generating workspace points efficiently. The visualization module can be built by either surface-based or volume-based algorithms. Each class of these algorithms is more suitable than the others for certain applications. The criteria selection module interacts with the user and tailors the most appropriate techniques, from the workspace point generation and the visualization modules, based on the application requirements.

Communicating Shared Resources: A Model For Distributed Real-Time Systems

Richard Gerber
MS-CIS-91-57
GRASP LAB 273

The timing behavior of a real-time system depends not only on delays due to process synchronization, but also on the availability of shared resources. Most current real-time models capture delays due to process synchronization; however, they abstract out resource-specific details by assuming idealistic operating environments. On the other hand, scheduling and resource allocation algorithms used for real-time systems ignore the effect of process synchronization except for simple precedence relations between processes. To bridge the gap between these two disciplines, we have developed a methodology called Communicating Shared Resources, or CSR. In this dissertation we describe our approach to the specification and verification of real-time systems. Application processes are specified in the CSR application language, which includes language constructs that are essential in real-time settings, such as timeouts, deadlines, periodic processes, interrupts and exception-handling. Then, a configuration schema is used to map the processes to system resources, and to specify the physical communication links between them. To analyze and execute the entire system, we automatically translate the result of the mapping into the CCSR process algebra. CCSR characterizes CSR's resource-based computation model by a priority-sensitive, operational semantics. To do this, we have formulated a natural treatment of preemption, which is based not only on priority, but also on resource utilization and inter-resource synchronization. The preemption ordering leads to a compositional proof system, which allows the syntactic manipulation of CCSR terms. Using this proof system, we perform the algebraic verification of our original real-time system.

Data Abstraction and General Recursion

Ramesh Subrahmanyam
MS-CIS-91-58
LOGIC & COMPUTATION 35

Existing approaches to semantics of algebraically specified data types such as Initial Algebra Semantics and Final Algebra Semantics do not take into account the possibility of general recursion and hence non-termination *in the ambient programming language*. Any technical development of this problem needs to be in the setting of domain theory. In this paper we present extensions of initial and final algebra semantics to algebras with an underlying domain structure. Four possibilities for specification methodologies arise: two each in the Initial and Final algebra paradigms. We demonstrate that the initial/final objects (as appropriate) exist in all four situations. The final part of the paper attempts to explicate the notion of abstractness of ADT's by defining a notion of operational semantics for ADT's, and then studying the relationship between the various algebraic-semantics proposed and the operational semantics.

Bounded Linear Logic

Jean-Yves Girard
Andre Scedrov
Philip J. Scott
MS-CIS-91-59
LOGIC & COMPUTATION 36

A typed, modular paradigm for polynomial time computation is proposed.

Equational and Rule-Based Programming: Visualization, Reliability, and Knowledge Base Generation

Jee-In Kim
MS-CIS-91-60

This document describes developing an environment for effective use of functional/equational programs and rule-based expert systems. There are significant advantages in using these paradigms for reliability, parallelism, and accumulation of expertise in knowledge bases. The environment will make it easier to understand and use these paradigms, construct more reliable systems, and automatically enrich rule-based knowledge bases with the expertise.

It will consist of the following components: (1) Visualization: for composing systems using a graphical interface and for understanding of algorithms. (2) Consistency Checking: for an equational and a rule-based languages in accordance with the semantics of the languages. (3) Knowledge Base Generation and Testing: a translator that extracts expertise from existing programs and accumulates it as rules in knowledge bases; the rules are tested to enhance reliability. (4) Verification: interactive heterogeneous reasoning that consists of equational reasoning based on visual and textual information.

These tools will be integrated in the proposed environment. The environment will greatly reduce the costs and increase the reliability of functional/equational and rule-based systems.

Forms of Semantic Specification

Carl A. Gunter

MS-CIS-91-61
LOGIC & COMPUTATION 37

The way to specify a programming language has been a topic of heated debate for some decades and at present there is no consensus on how this is best done. Real languages are almost always specified informally; nevertheless, precision is often enough lacking that more formal approaches could benefit both programmers and language implementors. My purpose is to look at a few of these formal approaches in hope of establishing some distinctions or at least stirring some discussion.

Surface Structure, Intonation, and "Focus"

Mark Steedman

MS-CIS-91-63
LINC LAB 206

The paper briefly reviews a theory of intonational prosody and its relation syntax, and to certain oppositions of discourse meaning that have variously been called "topic and comment", "theme and rheme", "given and new", or "presupposition and focus." The theory, which is based on Combinatory Categorial Grammar, is presented in full elsewhere. The present paper examines its implications for the semantics of "focus".

Towards Goal-Directed Diagnosis (Preliminary Report)

Ron Rymon
(University of Pennsylvania)
Bonnie L. Webber
(University of Pennsylvania)
John R. Clarke
(Medical College of Pennsylvania)

MS-CIS-91-67
LINC LAB 208

Recent research has abstracted diagnosis away from the *activity* needed to acquire information and to act on diagnosed disorders. In some problem domains, however, such abstraction is counter-productive and does not reflect real-life practice, which *integrates* diagnostic and therapeutic activity. Trauma management is a case in point. Here, we discuss a formalization of the integrated approach taken in TraumAID, a system we have developed to serve as an artificial aide to residents and physicians dealing with multiple trauma.

Among other things, the active pursuit of information raises the question of what is and what is not worth pursuing. In TraumAID 2.0, we take the view that the process of diagnosis should continue only as long as it is likely to make a difference to future actions. That view is formalized in the *goal-directed* diagnostic paradigm (GDD). Unlike other diagnostic paradigms, goal-directed diagnosis is first and foremost concerned with setting goals based on its conclusions. It regards the traditional construction of an explanation for the faulty behavior as secondary.

In order to explicitly represent goal-directedness, the diagnostic *process* is viewed as search in a space of attitude-beliefs. ¿From this, we derive a high-level algorithm that produces appropriate requests for action *while* searching for an explanation. A complete explanation, however, is not the criterion for terminating action. Such a criterion, we argue, is better treated in terms of goal-means tradeoffs. TraumAID's architecture, in so far as it embodies this goal-directed approach, assigns to a complementary *planner* the resolution of such tradeoffs.

Contact Operations Using An Instrumented Compliant Wrist

Thomas Lindsay
Janez Funda
Richard Paul

MS-CIS-91-68
GRASP LAB 275

Teleprogramming was developed as a solution to problems of teleoperation systems with significant time delays [5]. In teleprogramming, the human operator interacts in real

time with a graphical model of the remote site, which provides for real time visual and force feedback. The master system automatically generates symbolic commands based on the motions of the master arm and the manipulator/model interactions, given predefined criteria of what types of motions are to be expected. These commands are then sent via a communication link, which may delay the signals, to the remote site. Based upon a remote world model, predefined and possibly refined as more information is obtained, the slave carries out commanded operations in the remote world and decides whether each step has been executed correctly.

Contact operations involve the remote site manipulator interacting with the environment, including planned collisions, and motion with contact with the environment. A hybrid position/force control scheme using a instrumented compliant wrist has been demonstrated to be very effective for these types of operations. In particular, switching between position and force modes (when contacting a surface, for example) does not present problems for the system. A brief introduction of teleprogramming and contact operations is presented, including a model of sliding motions and early experimental results. Problems with these early experiments are presented, and solutions discussed. The criteria for an object to slide rather than tip over are presented, relating to the geometry of the object and the applied forces. Finally, methods are presented to match the experimental results to a simple model, to help the remote manipulator to quickly and robustly sense collisions.

Investigating Logics For Feasible Computation

Anuj Dawar

MS-CIS-91-69

Dynamics Of Rigid Bodies Undergoing Multiple Frictional Contacts

Yin-Tien Wang
Vijay Kumar
Jacob Abel

MS-CIS-91-70
GRASP LAB 276

There are several applications in robotics and manufacturing in which nominally rigid objects are subject to multiple frictional contacts with other objects. In most previous work, rigid body models have been used to analyze such systems. There are two fundamental problems with such an approach. Firstly, the use of frictional laws, such as Coulomb's law, introduce inconsistencies and ambiguities when used in conjunction with the principles of rigid body dynamics. Secondly, hypotheses traditionally used to model frictional impacts can lead to solutions which violate principles of energy conservation. In this paper these problems are explained with the help of examples. A new approach to the simulation of mechanical systems with multiple, frictional constraints is proposed which is free of inconsistencies.

Parallel Algorithms For Depth-First Search

Jon Freeman

MS-CIS-91-71

In this paper we examine parallel algorithms for performing a depth-first search (DFS) of a directed or undirected graph in sub-linear time. this subject is interesting in part because DFS seemed at first to be an inherently sequential process, and for a long time many researcher believed that no such algorithms existed. We survey three seminal papers on the subject. The first one proves that a special case of DFS is (in all likelihood) inherently sequential; the second shows that DFS for planar undirected graphs is in NC; and the third shows that DFS for general undirected graphs is in RNC. We also discuss randomnized algorithms, P-completeness and matching, three topics that are essential for understanding and appreciating the results in these papers

Abstract Syntax and Logic Programming

Dale Miller

MS-CIS-91-72
LINC LAB 209

When writing programs to manipulate structures such as algebraic expressions, logical formulas, proofs, and programs, it is highly desirable to take the linear, human-oriented, concrete syntax of these structures and parse them into a more computation-oriented syntax. For a wide variety of manipulations, concrete syntax contains too much useless information (*e.g.*, keywords and white space) while important information is not explicitly represented (em e.g., function-argument relations and the scope of operators). In parse trees, much of the semantically useless information is removed while other relationships, such as between function and argument, are made more explicit. Unfortunately, parse trees do not adequately address important notions of object-level syntax, such as bound and free object-variables, scopes, alphabetic changes of bound variables, and object-level substitution. I will argue here that the *abstract syntax* of such objects should be organized around α-equivalence classes of λ-terms instead of parse trees. Incorporating this notion of abstract syntax into programming languages is an interesting challenge. This paper briefly describes a logic programming language is presented to illustrate its approach to handling object-level syntax. A model theoretic semantics for this logic programming language is also presented.

Descriptive Complexity Approaches To Inductive Inference: A Critical Review
Kevin Atteson
MS-CIS-91-73
GRASP LAB 277

We present a general introduction and critical review of descriptive complexity approaches to inductive inference, that is, the problem of determining a model from observation. Descriptive complexity, as defined by Kolmogogov, Chaitin and Solomonoff is presented as a generalization of Shannons entropy. Its relations with randomness is discussed and examples are presented. the practicability of descriptive complexity theory is discussed. We then present Rissanen's MDL principle as a restriction of descriptive complexity. We demonstrate the effectiveness of MDL by applying it it AR processes. Lastly, we present and discuss LeClerc's application of MDL to image segmentation.

Constructive Logics Part I: A Tutorial On Proof Systems and Typed λ-Calculi
Jean Gallier
MS-CIS-91-74
LOGIC & COMPUTATION 40

The purpose of this paper is to give an exposition of material dealing with constructive logic, typed λ-calculi, and linear logic. The emergence in the past ten years of a coherent field of research often named "logic and computation" has had two major (and related) effects: firstly, it has rocked vigorously the world of mathematical logic; secondly, it has created a new computer science discipline, which spans from what is traditionally called theory of computation, to programming language design. Remarkably, this new body of work relies heavily on some "old" concepts found in mathematical logic, like natural deduction, sequent calculus, and λ-calculus (but often viewed in a different light), and also on some newer concepts. Thus, it may be quite a challenge to become initiated to this new body of work (but the situation is improving, there are now some excellent texts on this subject matter). This paper attempts to provide a coherent and hopefully "gentle" initiation to this new body of work. We have attempted to cover the basic material on natural deduction, sequent calculus, and typed λ-calculus, but also to provide an introduction to Girard's linear logic, one of the most exciting developments in logic these past five years. The first part of these notes gives an exposition of background material (with the exception of the Girard-translation of classical logic into intuitionistic logic, which is new). The second part is devoted to linear logic and proof nets.

Communication and Coarticulation In Facial Animation (Ph.D Disstertation)
Catherine Pelachaud
MS-CIS-91-77
GRAPHICS LAB 44

Our goal is to produce a high level programming language or tool for 3D animation of facial expressions, especially, those conveying information correlated with the intonation of the voice: this includes the differences of timing, pitch, and emphasis that are related to such semantic distinctions of discourse as "given" and "new" information, some of which are also correlated with affect or emotion. Up till now, systems have not embodies such rule-governed translation from speech and utterance meaning to facial expressions. Our algorithm embodies rules that describe and coordinate these relations (*intonation/information*, *intonation/emotions* and *facial expressions/emotions*). Given an utterance, we consider how the discourse information (what is *new/old* information in the given context, or what is the "topic" of the discourse) is transmitted through the choice of accents and their placement, how it is conveyed over facial expression and how the two are coordinated. The facial model integrates the action at several levels, including individual muscle, group of muscles, and eye- and head-motion, as well as the propagation of or interaction of these movements, especially coarticulation effects. This study offers a higher level of representation of facial actions by grouping them into specialized functions (lip shapes for phonemes, eyebrow and head motions as emphatic movements). The major "key phrases" of this work involves the integration of FACS (facial notational system derived by P. Ekman and W. Friesen), and the Action Units (muscle actions); it offers a solution to lip synchronization as well as it provides a repertory of the different types of facial expressions involved with speech; it considers speaker/listener interaction. This representation is used to drive an animation system linked to facial motion.

Control of Multiple Arm Systems With Rolling Constraints
Xiaoping Yun
Vijay Kumar
Nilanjan Sarkar
Eric Paljug
MS-CIS-91-79
GRASP LAB 278

When multiple arms are used to manipulate a large object, it is necessary to maintain and control contacts between the object and effector(s) on one or more arms. The contacts are characterized by holonomic as well as nonholonomic con-

straints. This paper addresses the control of mechanical systems subject to nonholonomic constraints, rolling constraints in particular. It has been shown that such a system is always controllable, but cannot be stabilized to a single equilibrium by smooth feedback. In this paper, we show that the system is not input-state linearizable though input-output linearization is possible with appropriate output equations. Further, if the system is position-controlled (i.e., the output equation is a functions of position variables only), it has a zero dynamics which is Lagrange stable but not asymptotically stable. We discuss the analysis and controller design for planar as well as spatial multi-arm systems and present results from computer simulations to demonstrate the theoretical results.

A New Range Finding Method Using A Varifocal Mirror
Chang Li
Xiaoping Yun
MS-CIS-91-80
GRASP LAB 279

A new range finding method is proposed in this paper which makes use of a varifocal mirror. The three-dimensional object space is first discretized into a sequence of spherical shells with a specially designed nonlinear vibrating varifocal mirror. These discrete spherical shell images are then recorded by a video camera. A deblurring algorithm is introduced in this paper which is used to remove the blurred components in the images. Different depth ranges can be obtained by controlling the vibration amplitude and the direct current component of the driving wave for the varifocal mirror. The depth accuracy is adjusted by varying the vibration period of the varifocal mirror. This range finding technique can be made real time by increasing the frame frequency of the camera.

Interactive Postural Control Of Articulated Geometric Figures (Dissertation)
Cary B. Phillips
MS-CIS-91-82
GRAPHICS LAB 45

Interactive postural control is the process of interactively pushing, poking, and twisting parts of an articulated geometric figure for the express purpose of getting it into a desired posture. Many motion algorithms and computer animation techniques generate motion sequences based on starting and ending postures for geometric figures, but few of these techniques address the fundamental problem of specifying these postures. The goal of this thesis is to develop a system that allows us to specify postures of animate geometric figures in ways that suggest how we interact with real people.

The emphasis of this thesis is on real-time interactive 3D manipulation. The elements of the interaction techniques form a powerful vocabulary for describing postures and postural adjustments. The vocabulary is not a spoken or written one; rather, it includes verbs acted out by the user through the movement of input devices.

There are three major components to this work. The first component is a real-time 3D direct manipulation technique that allows the user to intuitively translate and rotate "handles" on objects using only a three button mouse as input. The second component is an inverse kinematics algorithm that uses the notion of constraints, or desired geometric relationships, to control postures of articulated figures. The inverse kinematics formulation is well suited to highly redundant figures. The final component is the system of behaviors. Behaviors provide coordination between the parts of the figure, so that when one part of a figure moves, the body reacts as a whole. One of the most important behaviors, and the one requiring the most coordination, is balance. The behaviors magnify the effect of the basic manipulation commands so that relatively few invocations of the commands are necessary to accomplish a complex positioning task.

Force-Closure Grasps With Two Palms
José-Antonio N. Caraza
Xiaoping Yun
MS-CIS-91-83
GRASP LAB 280

This paper studies force-closure grasps of rigid objects by using two palms. The two palms are instrumented with tactile sensors capable of detecting the presence of contacts, and are assumed to be respectively installed on two robotic manipulators capable of motion and force control. Established in this paper is an existence condition under which the two palms form a force-closure grasp. The salient feature of this condition is that it does not require the information on the shape of the object and the contact locations. A configuration of the two palms in contact with the object satisfying this condition is called a force-closure grasp configuration (FCGC). Further an algorithm is developed to check the condition for FCGC in terms of the position and orientation of the palms.

A Multiagent System For Intelligent Material Handling
Ruzena Bajcsy
Richard Paul
Xiaoping Yun
Vijay Kumar
MS-CIS-91-84
GRASP LAB 281

The goal of our research is to investigate manipulation, mobility, sensing, control and coordination for *a multiagent robotic system* employed in the task of material handling, in an unstructured, indoor environment. In this research, manipulators, observers, vehicles, sensors, and human operator(2) are considered by be agents. Alternatively, an agent can be a general-purpose agent (for example, a six degree of freedom manipulator on a mobile platform with visual force, touch and position sensors). Possible applic;ations for such a system includes handling of waste and hazardous materials, decontamination of nuclear plants, and interfacing between special purpose material handling devices in warehouses.

The fundamental research problems that will be studied are *organization*, or the decomposition of the task into subtasks and configuring the multiple agents with appropriate human interaction, *exploration*, or the process of exploring geometric, material and other properties about the environment and other agents, and *coordination*, or the dynamic control of multiple agents for manipulation and transportation of objects to a desired destination.

Robotic Sensorimotor Learning In Continuous Domains

Marcos Salganicoff
Ruzena Bajcsy
MS-CIS-91-85
GRASP LAB 282

We propose that some aspects of task based learning in robotics can be approached using nativist and constructivist views on human sensorimotor development as a metaphor. We use findings in developmental psychology, neurophysiology, and machine perception to guide a robotic learning system's level of representation both for actions and for percepts. Visually driven grasping is chosen as the experimental task since it has general applicability and it has been extensively researched from several perspectives. An implementation of a robotic system with a dexterous three fingered hand, compliant instrumented wrist, arm and vision is used to test these ideas. Several sensorimotor primitives (vision segmentation and manipulatory reflexes) are implemented in this system and may be thought of as the "innate" perceptula and motor abilities of the system.

Applying empirical learning techniques to real situations brings up some important issues such as observation sparsity in high dimensional spaces, arbitrary underlying functional forms of the reinforcement distribution and robustness to noise in exemplars. The well established technique of nonparametric projection pursuit regression (PPR) is used to accomplish reinforcement learning by searching for *generalization directions* determining projections of high dimensional data sets which capture task invariants. Additionally, the learning process generally implies failures along the way. Therefore, the mechanics of the untrained robotic system must be able to tolerate grave mistakes during learning and not damage itself. We address this by the use of an instrumented compliant robot wrist which controls impact forces.

Visual Observation Of A Moving Agent

Tarek M. Sobh
Ruzena Bajcsy
MS-CIS-91-86
GRASP LAB 283

We address the problem of observating a moving agent. In particular, we propose a system for observing a manipulation process, where a robot hand manipulates an object. A discrete event dynamic systems (DEDS) frame work is developed for the hand/object interaction over time and a stabilizing observer is constructed. Low-level modules are developed for recognizing the "events" that causes state transitions within the dynamic manipulation system. The work examines closely the possibilities for errors, mistakes and uncertainties in the manipulation system, observer construction process and event identification mechanisms. The system utilizes different tracking techniques in order to observe and recognize the task in an *active, adaptive* and *goal-directed* manner.

Sensorimotor Learning Using Active Perception In Continuous Domains

Marcos Salganicoff
Ruzena Bajcsy
MS-CIS-91-87
GRASP LAB 284

We propose that some aspects of task based learning in robotics can be approached using nativist and constructivist views on human sensorimotor development as a metaphor. We use findings in developmental psychology, neurophysiology, and machine perception to guide a robotic learning system's level of representation both for actions and for percepts. Visually driven grasping is chosen as the experimental task since it has general applicability and it has been extensively researched from several perspectives. An implementation of a robotic system with a dexterous three fingered hand, compliant instrumented wrist, arm and vision is used to test these ideas. Several sensorimotor primitives (vision segmentation and manipulatory reflexes) are implemented in this system and may be though of as the "innate" perceptual and motor abilities of the system.

Applying empirical learning techniques to real situations brings up some important issues such as observation sparsity in high dimensional spaces, arbitrary underlying functional forms of the reinforcement distribution and robustness to

noise in exemplars. The well established technique of non-parametric projection pursuit regression (PPR) is used to accomplish reinforcement learning by searching for {em generalization directions} determining projections of high dimensional data sets which capture task invariants. Additionally, the learning process generally implies failures along the way. Therefore, the mechanics of the untrained robotic system must be able to tolerate grave mistakes during learning and not damage itself. WE address this by the use of an instrumented compliant robot wrist which controls impact forces.

[Accepted: AAAI Fall Symposium on Sensory Aspects of Robot Intelligence, Asilomar, CA, November 14-17 1991]

Important Considerations In Force Control With Applications To Multi-Arm Manipulation

Eric Paljug
Tom Sugar
Vijay Kumar
Xiaoping Yun
MS-CIS-91-88
GRASP LAB 287

This paper addresses force control in overconstrained dynamic systems with special emphasis on robot control and multiarm coordination. Previous approaches to force control are studied and many of these are shown to be unsuitable for dynamic force control. Practical and theoretical considerations for designing force control algorithms are discussed. Experimental and simulation results that validate the theoretical findings are presented for a single-degree-of-freedom pneumatic force controller. Finally the theoretical development of a two-arm manipulation system with an extended statespace formulation and a computer simulation of the system are presented to illustrate the application of the basic ideas to a more complicated system.

Robotic Manipulation Using A Behavioral Framework (Dissertation)

Sanjay Agrawal
MS-CIS-91-90
GRASP LAB 287

We endevour to build a robotic manipulation system that is capable of functioning in an unstructured environment. In order to provide such functionality, the system must be capable of reacting in a compliant manner to any stimuli in the environment at the motor level, and use the same information to modify its existing plan of action.

To build a general enough system that can handle a wide range of manipulatory tasks we build motion primitives both at the task level, which we term motor programs, and at the motor level which we term reflexes. Motor programs specify the behavior of the system at the motor level in terms of the desired stimuli response characteristic. Thus we introduce the concept of behavior to represent the sensorimotor relationship that can be defined at different levels of detail. More complex behaviors are manifested by coordinating the behavior of a number of motor systems.

Borrowing extensively from advances in the area of neurobiology and neuropsychology, we mimic the human motor control system by creating a hierarchical control system in which we incorporate the motor programs at the highest level. At the lower levels we breakdown the specified behavior into a constituent reflex actions for each of the manipulators in the system. Between the task level and the motor programs, and a sensorimotor level that reactively coordinates the behavior of the manipulators in the system to produce intelligent behavior.

The behavioral robotic framework is used both on a graphics simulator as well as an experimental testbed. The test bed consists of two robot arms, one mounted with an articulated mechanical hand instrumented with tactile sensors, and the other mounted with a mobile range scanner. We test, and extend the hypothesis that complex actions can be accomplished using simpler but coordinated reflexive motions.

Active and Exploratory Perception

Ruzena Bajcsy
Mario Campos
MS-CIS-91-91
GRASP LAB 288

The main goal of this paper is to show that there is a natural flow from active perception through exploration of perceptual learning. We have attempted to conceptualize the perceptual process of an organism that has the top-level task of surviving in an unknown environment. During this conceptualization process, four necessary ingredients have emerged for either artificial or biological organisms. First, the sensory apparatus and processing of the organism must be active and flexible. Second, the organism must have exploratory capabilities. Third, the organism must be selective in its data acquisition process. Fourth, the organism must be able to learn. In the section on learning, we have clearly delineated the difference between what must be innate and what must be learned. In order to test our theory, we present the system's architecture that follows from the perceptual task decomposition. The predictions of this theory are that an artificial system can explore and learn about its environment modulo its sensors, manipulators, end effectors and exploratory procedures/attribute extractors. It can describe its world with respect to the built-in alphabet, that

is the set of perceptual primitives.

Randomized Algorithms For Packet Routing On The Mesh

Sanguthevar Rajasekaren

MS-CIS-91-92
GRASP LAB 289

Packet routing is an important problem of parallel computing since a fast algorithm for packet routing will imply 1) fast inter-processor communication, and 2) fast algorithms for emulating ideal models like PRAMs on fixed connection machines.There are three different models of packet routing, namely 1) Store and forward, 2) Multipacket, and 3) Cut through. In this paper we provide a survey of the best known randomized algorithms for store and forward routing, $k-k$ routing, and cut through routing on the Mesh Connected Computers.

$k-k$ Routing, $k-k$ Sorting, and Cut Through Routing On The Mesh

Sanguthevar Rajasekaren

MS-CIS-91-93
GRASP LAB 290

In this paper we present randomized algorithms for $k-k$ routing, $k-k$ sorting, and cut through routing. The stated resource bounds hold with high probability. The algorithm for $k-k$ routing runs in $\lceil \frac{k}{2} \rceil n + o(kn)$ steps. We also show that $k-k$ sorting can be accomplished within $\lceil \frac{k}{2} \rceil n + n + o(kn)$ steps, and cut through routing can be done in $\frac{3}{4}kn + \frac{3}{2}n + o(kn)$ steps.The best known time bounds (prior to this paper) for all these three problems were $kn + o(kn)$. $\frac{kn}{2}$ is a known lower bound for all the three problems (which is the bisection bound), and hence our algorithms are very nearly optimal. All the above mentioned algorithms have optimal queue length, namely $k + o(k)$. These algorithms also extend to higher dimensional meshes.

An Active Observer

Ruzena Bajcsy

MS-CIS-91-95
GRASP LAB 295

In this paper we present a framework for research into the development of an **Active Observer**. The components of such an observer are the low and intermediate visual processing modules. Some of these modules have been adapted from the community and some have been investigated in the GRASP laboratory, most notably modules for the understanding of surface reflections via color and multiple views and for the segmentation of three dimensional images into first or second order surfaces via superquadric/parametric volumetric models. However the key problem in Active Observer research is the control structure of its behavior based on the task and situation. This control structure is modeled by a formalism called Discrete Events Dynamic Systems (DEDS).

Specification and Analysis Of Resource-Bound Real-Time Systems

Richard Gerber
University of Maryland
Insup Lee
University of Pennsylvania

MS-CIS-91-96
LOGIC & COMPUTATION 43
DISTRIBUTED SYSTEMS LAB 9

We describe a layered approach to the specification and verification of real-time systems. Application processes are specified in the CSR application language, which includes high-level language constructs such as timeouts, deadlines, periodic processes, interrupts and exception-handling. Then, a configuration scema is used to map the processes to system resources, and to specify the physical communication links between them. To analyze and execute the entire system, we automatically translate the result of the mapping into the CCSR process algebra. CCSR characterizes CSR's resource-based computatin model by a priority-sensitive, operational semantics, which yields a set of equivalence-preserving proof rules. Using this proof system, we perform the algebraic verification of our original real-time system.

Proving Memory Management Invariants For A Language Based On Linear Logic

Jawahar Chirimar
Carl A. Gunter
Jon G. Riecke

MS-CIS-91-98
LOGIC & COMPUTATION 45

We develop a tool for the rigorous formulation and proof of properties of runtime memory management for a sample programming language based on a linear type system. Two semantics are described, one at a level of observable results of computations and one that describes linear connectives in terms of memory-management primitives. The two semantics are proven equivalent and the memory-management model is proven to satisfy fundamental correctness criteria for reference counts.

Active Observer: A Discrete Event Dynamic System Model For Controlling An Observer Under Uncertainty (Ph.D Dissertation)

Tarek Sobh
MS-CIS-91-99
GRASP LAB 296

In this work we establish a framework for the general problem of observation, which may be applied to different kinds of visual tasks. We construct "intelligent" high-level control mechanisms for active visual recognition of different processes within a hybrid dynamic system. We address the problem of observing a manipulation process in order to illustrate the ideas and motive behind our framework. Active and autonomous tracking mechanisms are developed for the observer agent which is completely decoupled from the manipulation agent.

We use a discrete event dynamic system as a high-level structuring technique to model the manipulation system. The formulation utilizes the knowledge about the system and the different actions in order to solve the observer problem in an efficient, stable and practical manner. The model uses different tracking mechanisms so that the observer can "see" the workspace of the manipulating robot. An automaton is developed for the hand/object interaction over time and a stabilizing observer is constructed. Low-level modules are developed for recognizing the visual "events" that causes state transitions within the dynamic manipulation system in real time. A coarse quantization of the manipulation actions is used in order to attain an active, adaptive and goal-directed sensing mechanism. The formulation provides high-level symbolic interpretations of the scene under observation. The discrete event framework is augmented with mechanisms for recovering the continuous parametric evolution of the scene under observation and for asserting the state of the manipulation agent.

This work examines closely the possibilities for errors, mistakes and uncertainties in the manipulation system, observer construction process and event identification mechanisms. We identify and suggest techniques for modeling these uncertainties. Ambiguities are allowed to develop and are resolved after finite time. Error recovery mechanisms are also devised. The computed uncertainties are utilized for navigating the observer automaton state space, asserting state transitions and developing a suitable tracking strategy for the observer agent. The computed uncertainties are utilized for navigating the observer automaton state space, asserting state transition and developing a suitable tracking strategy for the observer agent. The approach used can be considered as a framework for a variety of visual tasks, as it lends itself to be a practical and feasible solution that uses existing information in a robust and modular fashion. Theoretical aspects and experimental results of the work support adopting this framework as a new basis for performing several task-oriented recognition, inspection and observation of visual phenomena.

Vision For Navigation Using Two Road Cues

Gareth D. Funka-Lea
MS-CIS-91-100
GRASP LAB 297

An autonomous vehicle must be able to find and maintain visual contact with a negotiable path before it. In outdoor environments this generally means locating a road in front of the vehicle. this masters thesis presents a strategy for road tracking that uses two road cues in order to maintain better visual contact with a road than could be achieved with either cue alone. Until recently, most autonomous vehicles relied on a single cue to find the road in front of them. Using multiple measurements from the image we can produce a more robust system. Also, using two cues to find the road, we have been able to study the ability of the system to recover when contradictory evidence exits concerning the location of the road in front of the vehicle.

The two cues to the road's location we use are the road's surface shading and its boundaries. The road's surface properties as captured in an image are modeled by bi-variate polynomial surface patches of up to second order. This is the first time that shading information has been used in vision for autonomous navigation. The road's boundaries, and other line-like features of a road, are modeled by line segments fit the image edges. With these two cues, we have a complementary description of an image - the surface patches describe the continuous portions of the image and the line segments describe the discontinuous portions of the image. The two cues also provide different information about the three dimensional environment in which the road and vehicle exist. The surface patches provide information about the material characteristics of the road and the illumination properties of the scene. The line segments representing the sides of a constant width road provide information about the relationship between the road and the vehicle in space in accordance with the laws of perspective viewing.

The surface patches and line segments that model the road in one image are used along with knowledge of the vehicle's motion to predict the appearance of the road in a subsequent image. At each image frame the models are updated to take into account new aspects of the road's appearance. The modeling states with the system segmenting under guidance from a human operator an image of the scene in front of the vehicle.

E-kernel On The IBM Victor V256 Multiprocessor–An Experimental Platform For Parallel Systems

Dennis G. Shea

MS-CIS-91-101
LINC LAB 211

This thesis presents the design of an experimental platform for parallel systems–E-kernel on the Victor V256 parallel system. E-kernel is an embedding kernel developed for the support of program mapping and network reconfiguration. E-kernel supports two levels of embedding: in the first level, the embedding of a network topology onto Victor's 2-d mesh network, and in the second level, the embedding of a task graph onto the system network. The first level corresponds to network reconfiguration (through software) and the second level to program mapping (placement of processes to processors through software).

Victor is a partitionable, message passing, experimental parallel system that has been designed and developed in the modular microsystems group of the parallel systems department at the IBM Thomas J. Watson Research Center in part to support this research. Key architectural features incorporated into Victor, such as the need to partitioning and monitoring, have been driven by this research, and solutions in part are a result of it. The development of E-kernel translates some of the many theoretical results in graph embeddings into practical tools for program mapping and network reconfiguration in a parallel system. With hardware and software system supports like these, users are able to design their programs according to the most natural task graph topologies, with the system attempting the communication optimization automatically on its network. Victor V256, together with E-kernel, provides a rich experimentation environment for parallel systems.

Teleprogramming: Remote Site Research Issues: (Dissertation Proposal)

Thomas Lindsay

MS-CIS-92-01
GRASP LAB 298

This document proposes the development of the remote site workcell for teleoperation with significant communication delays (on the order of one to 20 seconds). In these situations, direct teleoperation becomes difficult to impossible due to the delays in visual and force feedback. Teleprogramming has been developed in order to overcome this problem. In teleprogramming, the human operator interacts in real time with a graphical model of the remote site, which provides for real time visual and force feedback. The master arm and the manipulator/model interactions, given predefined criteria of what types of motions are to be expected. These commands are then sent via a communication link, which may delay the signals, to the remote site. Based upon a remote world models, predefined and possibly refined as more information is obtained, the slave carries out commanded operations in the remote world and decides whether each step has been executed correctly.

The remote site receives commands sent via the delayed communication link. These commands must be parsed and translated into the local robot control language, which includes insertion of dynamic parameters that are not generated the master system. The commands are then executed by the hybrid position/force controller, and the resulting motions monitored for errors.

This proposal addresses the following remote site issues: low level manipulator control using an instrumented compliant wrist for sensory feedback, higher level command execution implementing dynamic parameters, and remote manipulator tool usage and control.

Curved Path Walking

Hyeongseok Ko
Norman I. Badler

MS-CIS-92-02
GRAPHICS LAB 46

Research on biped locomotion has focused on sagittal plane walking in which the stepping path is a straight line. Because a walking path is often curved in a three dimensional environment, a 3D locomotion subsystem is required to provide general walking animation. In building a 3D locomotion subsystem, we tried to utilize pre-existing straight path (2D) systems. The movement of the center of the body is important in determining the amount of banking and turning. The center site is defined to be the midpoint between the two hip joints. An algorithm to obtain the center site trajectory that realizes the given curved walking path is presented. From the position and orientation of the center site, we compute stance and swing leg configuration as well as the upper body configuration, based on the underlying 2D system.

Interactive Image Display For The X Window System

John Bradley

MS-CIS-92-04
GRASP LAB 299

This report describes the program XV, which is an interactive color image display program for workstations and terminals running the X Window System. The program displays images saved in a variety of popular formats. It lets you arbitrarily stretch or compress the size of the image, rotate the image in 90-degree steps, flip the image around horizontal or vertical axes, crop off unwanted portions of the image, and measure pixel values and coordinates. Modified images can be saved in a variety of formats, or sent to a PostScript printer.

The program also features extensive color manipulation functions, including a colormap editor, hue remapping, brightness and contrast adjustment, and individual mapping functions for the Red, Green, and Blue video channels, to correct for device-dependent non-linear color response.

Multiple Instantiation Of Predicates In A Connectionist Rule-Based Reasoner

D.R. Mani
Lokendra Shastri

MS-CIS-92-05
LINC LAB 212

Shastri and Ajjanagadde have described a neurally plausible system for knowledge representation and reasoning that can represent systematic knowledge involving n-ary predicates and variables, and perform a broad class of reasoning with extreme efficiency. The system maintains and propagates variable bindings using temporally synchronous — i.e., in-phase — firing of appropriate nodes. This paper extends the reasoning system to incorporate *multiple instantiation of predicates*, so that any predicate can now be instantiated with up to k dynamic facts, k being a system constant. The ability to accommodate multiple instantiations of a predicate allows the system to handle a much broader class of rules; the system can even handle limited recursion (up to k levels). Though the time and space requirements increase by a constant factor, the extended system can still answer queries in time proportional to the length of the shortest derivation of the query and is independent of the size of the knowledge base.

Generating Human Motion By Symbolic Reasoning
Moon Ryul Jung
Norman I. Badler
MS-CIS-92-06
GRAPHICS LAB 47

This paper describes work on applying AI planning methods to generate human body motion for the purpose of animation. It is based on the fact that although we do not know how the body actually controls massively redundant degrees of freedom of its joints and moves in given situations, the appropriateness of specific behavior for particular conditions can be axiomatized at a gross level using commonsensical observations. Given the motion axioms (rules), the task of the planner is to find a discrete sequence of intermediate postures of the body via goal reduction reasoning based on the rules along with a procedure to discover specific collision-avoidance constraints, such that any two consecutive postures are related via primitive motions of the feet, the pelvis, the torso, the head, the hands, or other body parts. Our planner also takes account of the fact that body motions are continuous by taking advantage of execution-time feedback. Planning decisions are made in the *task space* where our elementary spatial intuition is preserved as far as possible, only dropping down to a *joint space formulation* typical in robot motion planning when absolutely necessary. We claim that our work is the first serious attempt to use an AI planning paradigm for animation of human body motion.

Goals and Actions In Natural Language Instructions (Dissertation Proposal)
Barbara Di Eugenio
MS-CIS-92-07
LINC LAB 213

Human agents are extremely flexible in dealing with Natural Language instructions: they are able both to adapt the plan they are developing to the input instructions, and vice versa, to adapt the input instructions to the plan they are developing. Borrowing the term from [Lewis 79], I call this two-way adaptation process accommodation.

In this proposal, I first define accommodation in the context of processing instructions. I then provide evidence for the particular inferences I advocate, and for the further claim that such inferences are directed by the goal to achieve which a certain action is performed. The evidence I provide comes from my analysis of naturally occurring instructions, and in particular of purpose clauses and of negative imperatives.

Finally, I propose a computational model of instructions able to support accommodation inferences. Such model is composed of: a speaker / hearer model of imperatives, based on the one presented in [Cohen and Levesque 90]; an action representation formalism based on a hybrid system, a' la KRYPTON [Brachman et al. 83], whose primitives are those proposed in [Jackendoff 90]; and inference mechanisms that contribute to building the structure of the intentions that the agent develops while interpreting instructions.

Collision-Free Path and Motion Planning For Anthropometric Figures
Wallace Ching
Norman I. Badler
MS-CIS-92-09
GRAPHICS LAB 48

This paper describes a collision free path planning and animation system for anthropometric figures. It can also take into consideration the strength limit of human figures and plan the motion accordingly. The algorithm breaks down the degrees of freedom of the figure into *Cspace groups* and computes the free motion for each of these groups in a sequential fashion. It traverses the tree in a depth first order to compute the motion for all the branches. A special playback routine is then used to traverse the tree in a reverse order to playback the final motion. Strength value measures are incorporated directly into the searching function so that path computed will obey strength availability criteria. The planner runs in linear time with respect to the total number of Cspace groups. The planner can interface with other simulation techniques to simulate complex human motions. We believe that the planner would find a path in most cases and is fast enough for practical use in a wide range of computer graphics applications.

The Reality of Virtual Environments
WPE II Paper
Rebecca T. Mercuri
MS-CIS-92-10
GRAPHICS LAB 49

Recent advances in computer technology have made it now possible to create and display three-dimensional virtual environments for real-time exploration and interaction by a user. This paper surveys some of the research done in this field at such places as: NASA's Ames Research Center, MIT's Media Laboratory, The University of North Carolina at Chapel Hill, and the University of New Brunswick. Limitations to the "reality" of these simulations will be examined, focusing on input and output devices, computational complexity, as well as tactile and visual feedback.

Superquadric Library, User Manual and Utility Programs
Luca Bogoni
MS-CIS-92-11
GRASP LAB 300

Superquadrics are a family of parametric shapes that have been used as primitives for shape representation in computer vision and computer graphics. They can be used for modeling tapering and bending deformations and are recovered efficiently by a stable numerical procedure.

This document introduces the superquadric library, *SQ_lib*, developed at the GRASP Lab at the University of Pennsylvania.

The manual is organized into three parts. The first part provides the reader with a description of superquadrics models and deformations that can be performed. Furthermore, it introduces the coordinate systems conventions which are used in the library.

The second part presents some examples of applications on how one can use the functions defined in the library. It also lists utility programs which have been developed while conducting research. They provide a good source of examples for the application of the library.

Finally, the last part describes the datatypes and each of the functions which are supported in the library. The library itself is organized in two sets *Fundamental* and *Auxiliary* functions. A quick reference to all the functions and an index is provided.

Some of the functions and examples supplied perform data preprocessing and are connected to the PM image description also available from the GRASP Lab. These functions are provided in isolation from the remaining body of the library and can easily be excluded in the actual compilation of the library. Furthermore, routines for the visualization of the data, using X11, are also provided.

Robotic Exploration Of Surfaces and Its Application To Legged Locomotion

Pramath Raj Sinha
MS-CIS-92-12
GRASP LAB 301

Material properties like penetrability, compliance, and surface roughness are important in the characterization of the environment. While concentrating on issues of geometry and shape, researchers in perceptual robotics, until recently, have not quite addressed the issue of the extraction of material properties from the environment. the goal of this research is to design and implement a robotic system that will actively explore a surface to extract its material characteristics. Further, the relevance of material properties in the legged locomotion of robots is also recognized and our research objectives are extended towards building a robotic system for exploration such that it actively perceives material properties during the process of legged locomotion. The chosen approach to the design and implementation of such a robotic system is to first select an appropriate environment model and then to design exploratory procedures salient to each attribute of interest. These exploratory procedures are then implemented through an experimental setup and the results show that material properties can be reliably measured. The design, implementation, and results of a framework for surface exploration to recover material properties are presented.

Further, the exploratory procedures for exploration are integrated into an active perceptual scheme for legged locomotion. The perceptual scheme is designed around creating the ability for the robot to sense variations in terrain properties while it is walking, so that is may be able to avoid sinking, slipping, and falling due to unexpected changes in the terrain properties, and make suitable changes in its foot forces to continue locomotion. Finite element simulations of the foot–terrain interaction are used to justify some of the strategies used in this active perceptula scheme. the active perceptual scheme is implemented by simulating a leg-ankle-foot system with a PUMA arm-compliant wrist-foot system and an accelerometer mounted on the foot to detect slip. Details of implementation and experimental results are presented.

Understanding Of Surface Reflections In Computer Vision By Color and Multiple Views (Dissertation)

Sang Wook Lee
MS-CIS-92-13
GRASP LAB 301

This thesis addresses problems and presents models for the detection and separation of specularities from Lambertian reflections using color and multiple images with different viewing direction. From the models, three algorithms are proposed and experimental results are presented. The first algorithms uses only color information for the separation of diffuse as well as sharp specularities and inter-reflections from Lambertian reflections through image segmentation. A computational model based on the dichromatic model is presented for interpretation of various surface reflections in a spectral space with three orthogonal basis functions. The established model is used for arranging color data for segmentation and separation. Applicable objects and illumination for the algorithm are limited to uniformly colored dielectrics under singly colored scene illumination. Use of multiple views for understanding reflection properties is proposed with the second and the third algorithms called spectral differencing and view sampling, respectively. Both use multiple views in different viewing directions, and are based on the Lambertian consistency that image irradiance from Lambertian reflection does not vary depending on viewing directions, while image irradiance from specular reflection or from the mixture of Lambertian and specular reflections

can change. Spectral differencing is a detection algorithm that detects specularities by color difference between two images without relying on any geometric feature correspondence. The object and illumination domain for detection is extended to nonuniformly colored dielectrics and metals under multiply colored scene illumination. With densely sampled views in wide angle and with known viewing directions, the view sampling algorithm reconstructs object structure as well as separates specularities from Lambertian reflections. The view sampling algorithm does not require color information, and is applicable to dielectrics and metals. Experimental results conform to the models and algorithms within the limitations discussed.

Planning Responses From High-Level Goals: Adopting The Respondent's Perspective In Cooperative Response Generation (Dissertation)
Brant A. Cheikes
MS-CIS-92-14
LINC LAB 215

Within the natural-language research community it has long been acknowledged that the conventions and pragmatics of natural-language communication often oblige dialogue systems to consider and address the underlying purposes of queries in their responses rather than answering them literally and without further comment or elaboration. Such systems cannot simply translate their users' requests into transactions on database or expert systems, but must apply many more complex reasoning mechanisms to the task of selecting responses that are both appropriate and useful. This idea has given rise to a broadly-defined program of research in *cooperative response generation* (CRG).

Research in CRG carried on over more than a decade has yielded a substantial body of literature. analysis of that literature, however, shows that investigators have focused primarily on modeling *manifestations* of cooperative behavior without directly considering the nature and motivations of the behavior itself. But if we want to develop natural-language dialogue systems that are truly to function as *cooperative respondents* instead of serving only as models of particular kinds of cooperative responses, a different approach is required.

I identify two opposing perspectives on the process of cooperative response generation: the **questioner-based** and the **respondent-based** perspectives. I argue that past research efforts have largely been questioner-based, and that this view has lead to the development of theories that are incompatible and cannot to integrated. I propose the respondent-based view as an alternative, and provide evidence that taking such a perspective might allow several interesting but otherwise poorly-understood aspects of cooperative response behavior to be modeled.

The final portion of the dissertation explores the computational implications of a respondent-based perspective. I outline the architecture of a **Cooperative Response Planning System**, a dialogue system that raises, reasons about, and attempts to satisfy high-level cooperative goals in its responses. This architecture constitutes a first approximation to a theory of how a system might reason from the beliefs it derives from a questioner's utterances to choose a cooperative response. The processing of two sample responses in this framework is described in detail to illustrate the architecture's capabilities.

Grasp Laboratory News
Volume 8, Number 1
Faculty & Graduate Students
MS-CIS-92-15
GRASP LAB 302

The GRASP laboratory has undertaken a research project to develop a multiagent robotic system for intelligent material handling. This article outlines the organization of the system. The goal of this research project is to investigate the architecture, coordination, and monitoring of a *multiagent robotic system* employed in the task of material handling, in an unstructured, indoor environment. In this research, manipulators, observers, vehicles, sensors, and human operators(s) are considered to be agents. Alternatively, an agent can be a general-purpose agent (for example, a six degree of freedom manipulator on a mobile platform and visual, force, touch and position sensors). Possible applications for such a system includes handling of waste and hazardous materials, decontamination of nuclear plants, and interfacing between special purpose material handling devices in warehouse. The fundamental research problems that are being studied are *organization*, or the decomposition of the task into subtasks and configuring the multiple agents with appropriate human interaction, *exploration*, or the process of exploring geometric, material and other properties about the environment and other agents, and coordination, or the dynamic control of multiple agents for manipulation and transportation of objects to a desired destination.

Simulation Of Mechanical Systems With Multiple Frictional Contacts
Yin-Tien Wang
Vijay Kumar
MS-CIS-92-16
GRASP LAB 303

There are several applications in robotics and manufacturing in which nominally rigid objects are subject to mul-

tiple frictional contacts with other objects. In most previous work, rigid body models have been used to analyze such systems. There are two fundamental problems with such an approach. Firstly, the use of frictional laws, such as Coulomb's law, introduce inconsistencies and ambiguities when used in conjunction with the principles of rigid body dynamics. Secondly, hypotheses traditionally used to model frictional impacts can lead to solutions which violate principles of energy conservation. In this paper these problems are explained with the help of examples. A new approach to the simulation of mechanical systems with multiple, frictional constraints is proposed which is free of inconsistencies.

Structural Recursion As A Query Language
Val-Breazu-Tannen
(University of Pennsylvania)
Peter Buneman
(University of Pennsylvania)
Shamim Naqvi
(BELLCORE)
MS-CIS-92-17
LOGIC & COMPUTATION 46

We propose a programming paradigm that tries to get close to both the semantic simplicity of relational algebra, and the expressive power of unrestricted programming languages. Its main computational engine is *structural recursion on sets*. All programming is done within a "nicely" typed lambda calculus, as in Machiavelli [OBB89]. A guiding principle is that *how* queries are implemented is as important as *whether* they can be implemented. As in relational algebra, the meaning of any relation transformer is guaranteed to be a total map taking finite relations to finite relations. A naturally restricted class of programs written with structural recursion has precisely the expressive power of the relational algebra. The same programming paradigm scales up, yielding query languages for the complex-object model [AB89]. Beyond that, there are, for example, efficient programs for transitive closure and we are also able to write programs that move out of sets, and then perhaps back to sets, as long as we stay within a (quite flexible) type system. The uniform paradigm of the language suggests positive expectations for the optimization problem. In fact, structural recursion yields *finer grain* programming therefore we expect that lower-level, and therefore better optimizations will be feasible.

Coordinating Locomotion and Manipulation Of A Mobile Manipulator
Yoshio Yamamoto
Xiaoping Yun
MS-CIS-92-18
GRASP LAB 304

A mobile manipulator in this study is a manipulator mounted on a mobile platform. Assuming the end point of the manipulator is guided, e.g., by a human operator to follow an arbitrary trajectory, it is desirable that the mobile platform is able to move as to position the manipulator in certain preferred configurations. Since the motion of the manipulator is unknown a priori, the platform has to use the measured joint position information of the manipulator for motion planning. This paper presents a planning and control algorithm for the platform so that the manipulator is always positioned at the preferred configurations measured by its manipulability. Simulation results are presented to illustrate the efficacy of the algorithm. The use of the resulting algorithm in a number of applications is also discussed.

Multi-Arm Manipulation Of Large Objects With Rolling Contacts
(Dissertation Proposal)
Eric D. Paljug
MS-CIS-92-19
GRASP LAB 305

We investigate the task of manipulating large objects relative to the size of the robot. In general, this task requires the use of more than one manipulator. The proposed approach is to use two robot arms, and to permit the robots to use any link surface for manipulation. Neither robot arm has a fixed grasp of the object. Its surface merely contacts the object surface. These contact points are capable of rolling, sliding and separation. The analysis begins by developing the equations that govern the system. Rolling at the contact points is included in the system model. Contact separation is avoid by enforcing the unilateral constraints that each arm must push at the contact point. Sliding is avoided by constraining the applied force to fall within the contact friction cone. A control algorithm is developed by employing nonlinear feedback obtained by applying techniques from differential geometry. The controller dynamically regulates force and motion simultaneously. A motion and force planner is also developed which incorporates the unilateral constraints into the system. The planner specifies a rolling motion for each contact which improves the system's ability to avoid slipping by repositioning the contact points such that forces are applied along the surface normals. The rolling motion is calculated based on the object dynamics and the desired critical contact force. Extensions to the theory are investigated by relaxing certain key assumptions. Simulations and experiments of the theoretical results are proposed to investigate the issues of practical implementation.

Proving Properties of Real-Time Distributed Systems: A Comparison of Three Approaches

Insup Lee

MS-CIS-92-20
GRASP LAB 306

Three formal methods for specifying properties of real-time systems are reviewed and used in a common example. Two of them offer a graphical representation and the third is an algebraic language. The example is that of an automatic railroad system with sensors to detect the train position and controls for the gate mechanism. Associated with each formalism is a proof methodology which is described and used to prove a safety property about the example. A comparison is made between the three formalisms according to various criteria including the expressiveness, readability, maintainability of the language, support for real-time concepts, method for expressing properties and proof mechanisms.

CLiFF Notes: Research In Natural Natural Processing at the University of Pennsylvania

Faculty & Graduate Students

MS-CIS-92-22
LINC LAB 217

The computational linguistics feedback forum (CLiFF) is a group of students and faculty who gather once a week to discuss the member's current research. As the work "feedback" suggest, the group's purpose is the sharing of ideas. The group also promotes interdiscipinary contacts between researchers who share an interest in Cognitive Science.

There is no single theme describing the research in Natural Language Processing at Penn. There is work done in CCG, Tree adjoining grammars, intonation, statistical methods, plan inference, instruction understanding, incremental interpretation, language acquisition, syntactic parsing, causal reasoning, free word order languages, and many other areas. With this in mind, rather than trying to summarize the varied work currently underway here at Penn, we suggest reading the following abstracts to see how the students and faculty themselves describe their work. Their abstract illustrate the diversity of interest amount the researchers, explain the areas of common interest, and describe some very interesting work in Cognitive Science.

This report is a collection of abstracts for both faculty and graduate students in Computer Science, Psychology and Linguistics. We pride ourselves on the close working relations between these groups, as we believe that the communication among the different departments and the ongoing interdepartmental research not only improves the quality of our work, but makes much of that work possible.

Progressive Horizon Planning–Planning Exploratory-Corrective Behavior

Ron Rymon
Bonnie L. Webber
John R. Clarke

MS-CIS-92-23
LINC LAB 218

Much planning research assumes that the goals for which one plans are known in advance. That is not true of trauma management, which involves both a search for relevant goals and reasoning about how to achieve them.

TraumAID is a consultation system for the diagnosis and treatment of multiple trauma. It has been under development jointly at the University of Pennsylvania and the Medical College of Pennsylvania for the past eight years. TraumAID integrates diagnostic reasoning, planning and action. Its reasoner identifies diagnostic and therapeutic goals appropriate to the physician's knowledge of the patient's state, while its planner advises on beneficial actions to next perform. The physician's lack of complete knowledge of the situation and the time limitations of emergency medicine constrain the ability of any planner to identify what would be the best thing to do. Nevertheless, TraumAID's *Progressive Horizon Planner* has been designed to create a plan for patient care that is in keeping with the standards of managing trauma.

Character Recognition Using A Modular Spatiotemporal Connectionist Model

Thomas Fontaine
Lokendra Shastri

MS-CIS-92-24
LINC LAB 219

We describe a connectionist model for recognizing handprinted characters. Instead of treating the input as a static signal, the image is scanned over time and converted into a time-varying signal. The temporalized image is processed by a spatiotemporal connectionist network suitable for dealing with time-varying signals. The resulting system offers several attractive features, including shift-invariance and inherent retention of local spatial relationships along the temporalized axis, a reduction in the number of free parameters, and the ability to process images of arbitrary length.

Connectionist networks were chosen as they offer learnability, rapid recognition, and attractive commercial possibilities. A modular and structured approach was taken in order to simplify network construction, optimization and analysis.

Results on the task of handprinted digit recognition are among the best report to date on a set of real-world ZIP code digit images, provided by the United States Postal Service.

The system achieved a 99.1% recognition rate on the training set and a 96.0% recognition rate on the test set with no rejections. A 99.0% recognition rate on the test set was achieved when 14.6% of the images.

Design, Implementation, and Evaluation Of A Real-Time Kernel For Distributed Robotics (Dissertation)

Robert Bruce King, II
MS-CIS-92-26
GRASP LAB 307

Modern robotics applications are becoming more complex due to greater numbers of sensors and actuators. The control of such systems may require multiple processor to meet the computational demands and to support the physical distribution of the sensors and actuators. A distributed real-time system is needed to perform the required communication and processing while meeting application-specified timing constraints. Our research is the design and evaluation of a real-time kernel, called *Timix V2*, for distributed robotics applications.

Timix V2 provides threads with dynamic timing constraints, execution environments as basic units for resource allocation and memory management context, and events to signal message arrival, device interrupts, alarms, and exceptions. The salient features of *Timix V2* are support for uniform scheduling and timely communication. *Timix V2* uses the notion of consistent scheduling to uniformly schedule both application and kernel threads to guarantee that the applications' real-time constraints are met. All device interrupt handlers, except the periodic clock interrupt, are converted to threads that are scheduled like any other thread. *Timix V2's* port-based message passing primitives support real-time communication by allowing individual message priorities to be used to order messages on a queue and by propagating scheduling information from a message to the associated thread on message arrival.

The kernel has been implemented on a distributed testbed and evaluated with respect to distributed real-time robotics applications.

A Robotic System for Learning Visually-Driven Grasp Planning (Dissertation Proposal)

Marcos Salganicoff
MS-CIS-92-27
GRASP LAB 308

We use findings in machine learning, developmental psychology, and neurophysiology to guide a robotic learning system's level of representation both for actions and for percepts. Visually-driven grasping is chosen as the experimental task since it has general applicability and it has been extensively researched from several perspectives. An implementation of a robotic system with a gripper, compliant instrumented wrist, arm and vision is used to test these ideas. Several sensorimotor primitives (vision segmentation and manipulatory reflexes) are implemented in this system and may be thought of as the "innate" perceptual and motor abilities of the system.

Applying empirical learning techniques to real situations brings up such important issues as observation sparsity in high-dimensional spaces, arbitrary underlying functional forms of the reinforcement distribution and robustness to noise in exemplars. The well-established technique of non-parametric projection pursuit regression (PPR) is used to accomplish reinforcement learning by searching for projections of high-dimensional data sets that capture task invariants.

We also pursue the following problem: how can we use human expertise and insight into grasping to train a system to select both appropriate hand preshapes and approaches for a wide variety of objects, and then have it verify and refine its skills through trial and error. To accomplish this learning we propose a new class of *Density Adaptive* reinforcement learning algorithms. These algorithms use statistical tests to identify possibly "interesting" regions of the attribute space in which the dynamics of the task change. They automatically concentrate the building of high resolution descriptions of the reinforcement in those areas, and build low resolution representations in regions that are either not populated in the given task or are highly uniform in outcome.

Additionally, the use of any learning process generally implies failures along the way. Therefore, the mechanics of the untrained robotic system must be able to tolerate mistakes during learning and not damage itself. We address this by the use of an instrumented, compliant robot wrist that controls impact forces.

Robust Location Estimation for MLR and Non-MLR Distributions (Dissertation Proposal)

Gerda L. Kamberova
MS-CIS-92-28
GRASP LAB 309

We address the problem of minimax location parameter estimation under zero-one loss. Consider the location data model $Z - \theta + V$. We observe the random variable Z and based on our observations(s), we want to estimate the value of the parameter θ, where we know a priori that $|\theta| \leq d$. The random noise V has a CDF F, F is independent of θ. The sampling distribution (the CDF of Z) is $F(z - \theta)$. The distribution f may be known, or it may be unknown.

In the latter case, there is a known class of distributions $\mathcal{F}$, to which F belongs. The locations data model is a starting point for considering more complex models, like $Z = h(\theta) + V$, where h is a nonlinear functions of θ, or more general $Z = h(\theta + V)$. The minimax criterion with zero-one loss is suitable for modeling problems in which it is desirable to minimize the maximum probability of getting unacceptable errors. Using this approach, as a consequence, we obtain fixed size confidence intervals, for the parameter, with highest probability of coverage.

There is a substantial difference in treating MLR and non-MLR distributions. When the sampling distribution is MLR, we obtain globally minimax decision procedures. They are monotone nondecreasing, continuous functions with very simple structures. When the sampling distribution is non-MLR, these monotone procedures are minimax within the class of all monotone decision rules. but they are not necessarily globally minimax. We explore the structure on non-monotone equalizer rules for non-MLR sampling distributions. We also explore the structure of Bayes rules. We believe that understanding of the structure of both equalizer and Bayes rules is a necessary step towards the delineation of global minimax decision procedures.

When the underlying distribution F is imprecisely known we consider the robust minimax decision problems. We study different sets of distributions (uncertainty classes) to which the CDF F may belong.

We use our results to address problems in sensor fusion. Many tasks in active perception require that we be able to combine different information from a variety of sensors which relate to one or more features of the environment. Prior to combining these data, we must test our observations for consistency. We examine sensor fusion problems for linear location data models. Our goal is to obtain: (i) a robust test of the hypothesis that data from different sensors are consistent; and (ii) a robust procedure for combining the data which pass this preliminary consistency test. Here, robustness refers to the statistical effectiveness of the decision rules when the probability distributions of the observation noise and the a priori position information associated with the individual sensors are uncertain.

Markov Random Field Models: A Bayesian Approach To Computer Vision Problems

Gerda L. Kamberova

MS-CIS-92-29
GRASP LAB 310

The object of our study is the Bayesian approach in solving computer vision problems. We examine in particular: (i) applications of *Markov random field* (MRF) models to modeling spatial images; (ii) MRF based statistical methods for image restoration, segmentation, texture modeling and integration of different visual cues.

Japanese Discourse and the Process of Centering

Marilyn Walker
(University of Pennsylvania)
Masayo Iida
(Hewlett Packard Laboratories)
Sharon Cote
(University of Pennsylvania)

MS-CIS-92-32
LINC LAB 220

This paper has two aims: (1) to generalize a computational account of discourse processing called CENTERING and apply it to discourse processing in Japanese, and (2) to provide some insights on the effect of syntactic factors in Japanese on discourse interpretation. We argue that while discourse interpretation is an inferential process, the syntactic cues constrain this process, and demonstrate this argument with respect to the interpretation of ZEROS, unexpressed arguments of the verb, in Japanese. The syntactic cues in Japanese discourse that we investigate are the morphological markers for grammatical TOPIC, the postposition *wa*, as well as those for grammatical functions such as SUBJECT, *ga* OBJECT, *o* and OBJECT2, *ni*. In addition, we investigate the role of speakers' EMPATHY, which is the perspective from which an event is described. This is morphologically indicated through the use of verbal compounding, i.e. that auxiliary use of verbs such as *kureta, kita*. Our results are based on a survey of native speakers of their interpretation of short discourses, consisting of minimal pairs, varied by one of the above factors. We demonstrate that these syntactic cues do indeed affect the interpretation of ZEROS, but that having previously been the TOPIC and being realized as a ZERO also contribute to an entity being interpreted as the TOPIC. We propose a new notion of TOPIC AMBIGUITY, and show that CENTERING provides constraints on when a ZERO can be interpreted as the TOPIC.

Conceptual Structures and CCG: Linking Theory and Incorporated Argument Adjuncts

Michael White

MS-CIS-92-34
LINC LAB 222

In Combinatory Categorial Grammar (CCG) [Steedman90,SteedmanLanguage91], semantic function-argument structures are compositionally produced through the course of a derivation. These structures identify, *inter alia*, which

entities play the same roles in different events for expressions involving a wide range of coordinate constructs. This sameness of role (i.e. *thematic*) information is not identified, however, across cases of verbal diathesis. To handle these cases as well, the present paper demonstrates how to adapt the solution developed in Conceptual Semantics [Jackendoff90,Jackendoff91] to fit the CCG paradigm.

The essence of the approach is to redefine the Linking Theory component of Conceptual Semantics in terms of CCG categories, so that derivations yield conceptual structures representing the desired thematic information; in this way no changes are required on the CCG side. While this redefinition is largely straightforward, an interesting problem arises in the case of Conceptual Semantics' Incorporated Argument Adjuncts. In examining these, the paper shows that they cannot be treated as adjuncts in the CCG sense without introducing new machinery, nor without compromising the independence of the two theories. For this reason, the paper instead adopts the more traditional approach of treating them as oblique arguments.

Control of Discrete Event Systems
Jana Košecká
MS-CIS-92-35
GRASP LAB 313

Discrete Event Systems (DES) are a special type of dynamic systems. The "state" of these systems changes only at discrete instants of time and the term "event" is used to represent the occurrence of discontinuous changes (at possibly unknown intervals). Different Discrete Event Systems models are currently used for specification, verification, synthesis as well as for analysis and evaluation of different qualitative and quantitative properties of existing physical systems.

The main focus of this paper is the presentation of the automata and formal language model for DES introduced by Ramadge and Wonham in 1985. This model is suitable for the examination of some important control theoretic issues, such as controllability and observability from the qualitative point of view, and provides a good basis for modular synthesis of controllers. We will also discuss an Extended State Machine and Real-Time Temporal Logic model introduced by Ostroff and Wonham in [Ostroff-ESM87]. It incorporates an explicit notion of time and means for specification and verification of discrete event systems using a temporal logic approach. An attempt is made to compare this model of DES with other ones.

Randomized Routing and Sorting On The Reconfigurable Mesh
Sanguthevar Rajasekaran
Theodore McKendall
MS-CIS-92-36
GRASP LAB 314

In this paper we demonstrate the power of reconfiguration by presenting efficient randomized algorithms for both packet routing and sorting on a reconfigurable mesh connected computer (referred to simply as the mesh from hereon). The run times of these algorithms are better than the best achievable time bounds on a conventional mesh.

In particular, we show that permutation routing problem can be solved on a linear array of size n in $\frac{3}{4}n$ steps, whereas $n-1$ is the best possible run time without reconfiguration. We also show that permutation routing on an $n \times n$ reconfigurable mesh can be done in time $n+o(n)$. In contrast, $2n-2$ is the diameter of a conventional mesh and hence routing and sorting will need at least $2n-2$ steps on a conventional mesh. In addition we show that the problem of sorting can be solved in time $n+o(n)$. All these time bounds hold with high probability. The bisection lower bound for both sorting and routing on the mesh is $\frac{n}{2}$, and hence our algorithms have nearly optimal time bounds.

An Active Approach To Functionality Characterization and Recognition
Luca Bogoni
Ruzena Bajcsy
MS-CIS-92-37
GRASP LAB 315

In this paper we focus on understanding and defining a methodology for object description and recognition both in terms of its geometrical, material and functional specifications. We define functionality in an object as its applicability toward the achievement of a task. We emphasize and develop an interactive and performatory approach to functionality recovery. Furthermore, we introduce the distinction between *Inherent, Intended* and *Imposed* functionality.

By analyzing interaction and manipulation tasks as goal–oriented recognition processes we propose to identify and characterize functionalities of objects. This interaction is not only a means of verification of the hypothesized presence of functionality in objects but also a way to actively and purposively recognize the object.

In order to accomplish our goal, we introduce a formal model, based on Discrete Event Dynamic System Theory, to define a task for recovering and describing functionality. We extend the recovery process to an algebra of tasks. We describe how a more complex task can be composed from

a set of primitive ones. This constructive approach allows a task to be built from simpler ones in an stepwise fashion.

Once the manipulatory task has been described in the formal model, it must be instantiated in a context. In such a context, the behavior of the system in which the interaction between a Manipulator, a Tool and a Target object must be observed. Thus, the description of tasks themselves provide must for means of addressing observability through different sensor modalities. For this purpose, we introduce the notion of *Partial Observability* of a task. This allows the description of a plant in which not all events and the time of their occurrence might be modelled and therefore predictable in advance.

Robotic Exploration Of Material and Kinematic Properties Of Objects (Dissertation)

Mario Fernando Montenegro Campos

MS-CIS-92-38

GRASP LAB 316

The physical interaction with unstructured environments requires that robotic systems have the ability to extract material and kinematic properties of objects around them. The goal of this research is to design a robotic systems that actively explores and extracts the material properties, including thermal, hardness and mass properties and of kinematic properties, such as mobility and geometric parameters of objects and their parts. to accomplish this objective, we invoke the paradigms of active perception and exploratory procedures. We develop methodologies for the design of such procedures as well as sensors which support their use in the robotic domain and demonstrate their effectiveness. The system is composed of a control module which coordinates the visual and the haptic sub-modules. Vision is implemented via an agile laser range-scanner which is able to acquire different views of the desired object. Global volumetric models of the object are recovered by fitting superellipsoids to the $2\frac{1}{2}$D range image. The haptic module uses the geometric information of the object to perform several tests based on non-destructive techniques. For exploring thermal properties, a new approach for the design and modeling of thermal sensors for robotics is presented. A model of this sensor is developed and its validity is experimentally verified with different materials. Mass density is estimated by the weight evaluation procedure. Hardness is evaluated by means of stress *vs.* strain tests. Compression and tension tests are performed to determine this property. The kinematic characteristics of the object are explored by the mobility procedure. We describe a novel methodology, based on screw theoretic results which enables the identification of the mobility of the object. This is accomplished by forming a closed kinematic chain with the manipulator and the unknown object. The number of degrees of freedom present in the object as well as the geometric parameters of its links are then extracted. The design and implementation of the robotic haptic architecture testbed where all of the above concepts were smoothly integrated into a working system is also described. The architecture controls and coordinates the two robot manipulators, the instrumented parallel-jaw gripper and the mobile laser range-scanner.

Parallel Evidence-Based Indexing of Complex Three-Dimensional Models Using Prototypical Parts and Relations (Dissertation Proposal)

Ron Katriel

MS-CIS-92-39

GRASP LAB 317

LINC LAB 223

This proposal is concerned with three-dimensional object recognition from range data using superquadric primitives. Superquadrics are a family of parametric shape models which represent objects at the part level and can account for a wide variety of natural and man-made forms. We propose a vision architecture that scales well as the size of its model database grows. Following the recovery of superquadric primitives from the input depth map, we split the computation into two concurrent processing streams. One is concerned with the classification of individual parts using viewpoint-invariant shape information while the other classifies pairwise part relationships using their relative size, orientation and type of joint. The major contribution of this proposal lies in a principled solution to the very difficult problems of superquadric part classification and model indexing. The problem is how to retrieve the best matched models without exploring all possible object matches. Our approach is to cluster together similar model parts to create a reasonable number of prototypical part classes (protoparts). Each superquadric part recovered from the input is paired with the best matching protopart using precomputed class statistics. A parallel, theoretically-well grounded evidential recognition algorithm quickly selects models consistent with the classified parts. Classified part relations (protorelations) are used to further reduce the number of consistent models and remaining ambiguities are resolved using sequential top-down search.

Occlusions As A Guide For Planning The Next View

Jasna Maver

Ruzena Bajcsy

MS-CIS-92-40

GRASP LAB 318

To resolve the ambiguities that are caused by occlusions in images, we need to take sensor measurements from several different views. The task addressed in this paper deals with a strategy for acquiring 3-D data of an unknown scene. We must first answer the question: What knowledge is adequate to perform a specific task? Thinking in the spirit of purposive vision, to accomplish its task, a system does not need to understand the complete scene but must be able to recognize patterns and situations that are necessary for accomplishing the task.

We have limited ourselves to range images obtained by a light stripe range finder. A priori knowledge given to the system is the knowledge of the sensor geometry. The foci of attention are occluded regions, i.e., only the scene at the borders of the occlusions is modeled to compute the next move. Since the system has knowledge of the sensor geometry, it can resolve the appearance of occlusions by analyzing them.

The problem of 3-D data acquisition is divided in two subproblems due to two types of occlusions. An occlusion arises either when the reflected laser light does not reach the camera or when the directed laser light does not reach the scene surface. After taking the range image of a scene the regions of no data due to the first kind of occlusion are extracted. The missing data are acquired by rotating the sensor system in the scanning plane, which is defined by the first scan. After a complete image of the surface illuminated from the first scanning plane has been built, the regions of missing data which are due to the second kind of occlusions are located. Then the directions of the next scanning planes for further 3-D data acquisition are computed.

For more detailed proofs and theorems of the computation of the next scanning plane (Section 5.3) please see "Occlusions as a Guide for Planning the Next View," University of Pennsylvania Technical Report, MS-CIS-91-27, GRASP LAB 257, 1991.

Analysis and Simulation of Mechanical Systems with Multiple Frictional Contacts (Dissertation)

Yin-Tien Wang
MS-CIS-92-41
GRASP LAB 319

There are several applications in robotics and manufacturing in which nominally rigid objects are subject to multiple frictional contacts. Since such a system is characterized by unilateral constraints, the topology of the mechanical system varies with time. that is, each time when a contact is formed or broken, or when a rolling contact changes to a sliding contact, the mobility of the mechanical system and the structure of the differential equations that characterize the system change. The research in this dissertation focuses on a systematic method for the analysis and simulation of such systems.

In most previous work, rigid body models and empirical models for friction have been used to analyze the dynamics of such systems. It is shown here that the use of frictional laws, such a Coulomb's law, introduce inconsistencies and ambiguities when used in conjunction with the principles of rigid body dynamics. Further, the static indeterminacy makes it impossible to determine the contact forces.

A new approach to the simulation of mechanical systems with multiple, frictional constraints is proposed which is free of inconsistencies. Compliant contact models are used to model the deformation of the contact surface and the energy dissipation during impacts. The method involves the integration of rigid body models with the compliant contact models - the rigid body models are used for predicting gross motion in the absence of unilateral constraints, and the contact models are used, when frictional contacts occur, for analyzing small motions. This method is compared with previous hypotheses and models and is shown to overcome their limitations.

The general method developed in this dissertation has applications in a wide range of problems in manufacturing and robotics. In this dissertation, we address the dynamic analysis and simulation of nonlinear control algorithms for multiarm manipulation, control of enveloping grasps and the parts-feeding processing manufacturing.

Pronominal Reference To Events and Actions: Computational Foundations (Dissertation)

Ethel Schuster
MS-CIS-92-43
LINC LAB 224

When a pronoun appears in discourse, it can refer to a specific event, to various types of events, as well as to sets of events. It is not always possible to identify a one-to-one correspondence between the pronoun and its referent. This thesis presents an approach whereby such a correspondence can be identified. Two types of relationships among referents are identified: (i) a generalization relationship, which establishes the relationship between a specific event, described in the discourse, and a general class of events, and (ii) three *compounding* relationships, *sequence, causation,* and *generation.* These compounding relationships connect various events as compound unit as a whole or to parts of it, depending on the particular compounding relationships that hold within the compound.

This thesis also presents a set of rules that guide the choice of the referents of the pronouns *it* and *that.* This

set of rules leads to an algorithm that generates pronouns referring to individual or compound events. By using one pronoun over the other, it is possible to indicate whether the pronoun refers to a compound referent or to parts of that compound.

Control of Mechanical Systems with Rolling Constraints: Application To Dynamic Control of Mobile Robots
Nilanjan Sarkar
Xiaoping Yun
Vijay Kumar
MS-CIS-92-44
GRASP LAB 320

There are many examples of mechanical systems which require rolling contacts between two or more rigid bodies. Rolling contacts engender nonholonomic constraints in an otherwise holonomic system. In this paper, we develop a unified approach to the control of mechanical systems subject to both holonomic and nonholonomic constraints. We first present a state space realization of a constrained system and show that it is not input-state linearizable. We then discuss the input-output linearization and zero dynamics of the system. This approach is applied to the dynamic control of mobile robots. Two types of control algorithms for mobile robots are investigated: (a) trajectory tracking, and (b) path following. In each case, a smooth nonlinear feedback is obtained to achieve asymptotical input-output stability, and Lagrange stability of the overall system. Simulation results are presented to demonstrate the effectiveness of the control algorithms and to compare the performance of trajectory tracking and path following algorithms.

On Feedback Linearization of Mobile Robots
Xiaoping Yun
Yoshio Yamamoto
MS-CIS-92-45
GRASP LAB 321

A wheeled mobile robot is subject to both holonomic and nonholonomic constraints. Representing the motion and constraint equations in the state space, this paper studies the feedback linearization of the dynamic system of a wheeled mobile robot. The main results of the paper are: (1) It is shown that the system is not input-state linearizable. (2) If the coordinates of a point on the wheel axis are taken as the output equation, the system is not input-output linearizable by using a static state feedback; (3) but is input-output linearizable by using a dynamic state feedback. (4) If the coordinates of a reference point in front of the mobile robot are chosen as the output equation, the system is input-output linearizable by using a static state feedback. (5) The internal motion of the mobile robot when the reference point moves forward is asymptotically stable whereas the internal motion when the reference point moves backward is unstable. A nonlinear feedback is derived for each case where the feedback linearization is possible.

From Operational Semantics To Abstract Machines
John Hannan
(University of Copenhagen)
Dale Miller
(University of Pennsylvania)
MS-CIS-92-46
LINC LAB 225

We consider the problem of mechanically constructing abstract machines from operational semantics, producing intermediate-level specifications of evaluators guaranteed to be correct with respect to the operational semantics. We construct these machines by repeatedly applying correctness-preserving transformations to operational semantics until the resulting specifications have the form of abstract machines. Though not automatable in general, this approach to constructing machine implementations can be mechanized, providing machine-verified correctness proofs. As examples we present the transformation of specifications for both call-by-name and call-by-value evaluation of the untyped λ-calculus into abstract machines that implement such evaluation strategies. We also present extensions to the call-by-value machine for a language containing constructs for recursion, conditionals, concrete data types, and built-in functions. In all cases, the correctness of the derived abstract machines follows from the (generally transparent) correctness of the initial operational semantic specification and the correctness of the transformations applied.

To appear in the Journal of Mathematical Structures in Computer Science.

Naturally Embedded Query Languages
Val Breazu-Tannen
Peter Buneman
Limsoon Wong
MS-CIS-92-47
LOGIC & COMPUTATION 47

We investigate the properties of a simple programming language whose main computational engine is *structural recursion* on sets. We describe a progression of sublanguages in this paradigm that (1) have increasing expressive power, and (2) illustrate robust conceptual restrictions thus exhibiting interesting additional properties. These properties suggest that we consider our sublanguages as candidates of

"query languages". Viewing query languages as restrictions of our more general programming language has several advantages. First, there is no "impedance mismatch" problem; the query languages are already there, so they share common semantic foundation with the general language. Second, we suggest a uniform characterization of nested relational and complex-object algebras in terms of some surprisingly simple operators; and we can make comparisons of expressiveness in a general framework. Third, we exhibit differences in expressive power that are not always based on complexity arguments, but use the idea that a query in one language may not be *polymorphically* expressible in another. Fourth, ideas of category theory can be profitably used to organize semantics and syntax, in particular our minimal (core) language is a well-understood categorical construction: a cartesian category with a strong monad on it. Finally, we bring out an *algebraic* perspective, that is, our languages come with equational theories, and categorical ideas can be used to derive a number of rather general identities that may serve as optimizations or as techniques for discovering optimizations.

π-Calculus As A Theory In Linear Logic: Preliminary Results
Dale Miller
MS-CIS-92-48
LINC LAB 226

The agent expressions of the π-calculus can be translated into a theory of linear logic in such a way that the reflective and transitive closure of π-calculus (unlabeled) reduction is identified with "entailed-by". Under this translation, parallel composition is mapped to the multiplicative disjunct ("par") and restriction is mapped to universal quantification. Prefixing, non-deterministic choice (+), replication (!), and the match guard are all represented using non-logical constants, which are specified using a simple form of axiom, called here a *process clause*. These process clauses resemble Horn clauses except that they may have multiple conclusions; that is, their heads may be the par of atomic formulas. Such multiple conclusion clauses are used to axiomize communications among agents. Given this translation, it is nature to ask to what extent proof theory can be used to understand the meta-theory of the π-calculus. We present some preliminary results along this line for π_0, the "propositional" fragment of the π-calculus, which lacks restriction and value passing (π_0 is a subset of CCS). Using ideas from proof-theory, we introduce *co-agent* and show that they can specify some testing equivalences for π_0. If negation-as-failure-to-prove is permitted as a co-agent combinator, then testing equivalence based on co-agents yields observational equivalence for π_0. This latter result follows from observing that co-agents are isomorphic to formulas in the Hennessy-Milner modal logic.

A Structural Interpretation of Combinatory Combinatory Categorial Grammar
James Henderson
MS-CIS-92-49
LINC LAB 227

This paper gives an interpretation of Combinatory Categorial Grammar derivations in terms of the construction of traditional phrase structure trees. This structural level of representation not only shows how CCG is related to other grammatical investigations, but this paper also uses it to extend CCG in ways which are useful for analyzing and parsing natural language, including a better analysis of coordination.

A Reconsideration of Preconditions
Christopher W. Geib
MS-CIS-92-50
LINC LAB 228

This paper is part of an attempt to introduce intentionality of the actor to planning decisions. As a first step in this process the usual representations for actions used by planning systems must be reevaluated. this paper argues for the elimination of preconditions and qualification conditions from action representation in favor of explicit representation of intention, situated reasoning about the results of the action and reactive failure mechanisms. The paper then describes a planning system that has explicit representation and use of intentions and uses action representation that do not have preconditions.

Surface Structure
Mark Steedman
MS-CIS-92-51
LINC LAB 229

The purpose of this paper is to show how binding and control can be captured straightforwardly in Combinatory Categorial Grammar (CCG), and to examine the interaction of the binding theory with the CCG account of long-range dependencies including "parasitic gaps" (Steedman 1987, Szabolcsi 1987a). Part I shows that a simple theory of binding and control is compatible with CCG. Part II shows that the Binding Theory interacts correctly with the combinatory account of long range dependency, correctly imposing certain constraints, including a number of asymmetries with respect to extraction between subjects and other arguments, such as "strong crossover", and the equivalent of an "anti-c-command" restriction on parasitic gaps (cf. Taraldsen). The conclusion suggests a simplifying reorganisation of the theory of grammar, via a single level of derivational syntax and a reallocation of responsibilities among the modules of the theory.

Polymorphism and Inference In Database Programming

Peter Buneman
(University of Pennsylvania)
Atsushi Ohori
(Oki Electric)
MS-CIS-92-72
LOGIC & COMPUTATION 52

The polymorphic typee system of ML can be extended in two ways to make it the appropriate basis of a database programming language. The first is an extension to the language of types that captures the polymorphic nature of field selection; the second is a technique that generalizes relational operators to arbitrary data structures. The combination provides a statically typed language in which relation database may be cleanly represented as typed structures. As in ML types are inferred, which relieves the programmer of making the rather complicated type assertions that may be required to express the most general type of program that involving field selection and generalized relational operators.

These extensions may also be used to provide static polymorphic typechecking in object-oriented languages and databases. A problem that arises with object oriented databses is the apprarent need for the dynamic typechecking when dealing with queries on heterogeneous collections of objects. An extension of the type system needed for generalized relational operations can also be used for manipulating collections of dynamically typed values in a statically typed language. A prototype language based on these ideas has been implemented. While it lacks a proper treatment of persistent data, it demonstrates that a wide variety of database structures can be cleanly represented in a polymorphic programming language.

Intentions In Means-End Planning (Dissertation Proposal)

Christopher W. Geib
MS-CIS-92-73
LINC LAB 235

This proposal discusses the use of the intentions of the actor in performing means-end reasoning. In doing so, it will show that preconditions and applicability conditions in existing systems are ill-defined and intrinscially encode situaltional information that prevents intentions from playing a role in the planning process. While the former problem can be fixed, the latter cannot. Therefore, I argue that preconditions should be eliminated from action representation. In their place, I suggest explicit representation of intention, situalted reasoning about the results of action, and robust failure mechanisms. I then describe a system, the Intentional Planning System (ItPlanS), which embodies these ideas, compare ItPlanS to other systems, and propose future directions for this work.

Doing What You're Told: Following Task Instruction In Changing, but Hospitable Environments

Bonnie Webber, Norman Badler,
F. Breckenridge Baldwin, Welton Becket
Barbara Di Eugenio, Christopher Geib
Moon Jung, Libby Levison
Michael Moore, Michael White
MS-CIS-92-74
LINC LAB 236

The AnimNL project (*Anim*ation from *N*atural *L*anguage) has as its goal the automatic creation of *animated task simulations* from *natural-language instructions*. The question addressed in this paper is how agents can perform tasks in environments about which they have only partial relevant knowledge. The solution we describe involves enabling such agents to

- develop expectations through instruction understanding and plan inference, and use those expectations in deciding how to act;
- exploit generalized abilities in order to deal with novel geometric situations.

The AnimNL project builds on an animation system, *Jack*™, that has been developed at the Computer Graphics Research Lab at the University of Pennsylvania, and draws upon a range of recent work in Natural Language semantics, planning and plan inference, philosophical studies of intention, reasoning about knowledge and action, and subsumption architectures for autonomous agents.

Model Based Teleoperation To Eliminate Feedback Delay NSF Grant BCS89-01352 Second Report

Richard P. Paul
Janez Funda
Thomas Lindsay
Masahiko Hashimoto
Craig Sayers
MS-CIS-92-75
GRASP LAB 333

We are conducting research in the area of teleoperation with feedback delay. Delay occurs with earth-based teleoperation in space and with surface-based teleoperation with untethered submersibles when acoustic communication links are involved. The delay in obtaining position and force feedback from remote slave arms makes teleoperation extremely difficult leading to very low productivity.

We have combined computer graphics with manipulator programming to provide a solution to the problem. A teleoperator master arm is interfaced to a graphics based simulator of the remote environment. The system is then coupled with a robot manipulator at the remote, delayed site. The operator's actions are monitored to provide both kinesthetic and visual feedback and to generate symbolic motion commands to the remote slave. The slave robot then executes these symbolic commands delayed in time. While much of a task proceeds error free, when an error does occur, the slave system transmits data back to the master environment which is then "reset" to the error state from which the operator continues the task.

Acquisition Of Verb Categories

Mark Steedman

MS-CIS-92-76

LINC LAB 237

The paper was delivered as a commentary upon Michael Brent's presentation "Acquisition of Subcategorization Frames Using Aggregated Evidence from Local Syntactic Cues" to the Conference on Acquisition of the Mental Lexicon, IRCS, University of Pennsylvania, January 1992. It argues in support of using statistical techniques like Brent's to minimise the consequences of errors and misanalyses, but concludes that the case for believing that children acquire subcategorisations and other aspects of syntax on the basis of semantic and contextual cues remains strong.

Improved Instrumented Compliant Wrist Design

Thomas Lindsay
Richard P. Paul

MS-CIS-92-77

GRASP LAB 334

Interaction between robot and environment is an extremely important aspect of robotic research. Compliance helps reduce the impact effects of robot/environment interaction. Hybrid position/force control is important in most robotic tasks; accurate position control is needed in unconstrained directioins, and accurate fore control is needed in constrained directions. Force control can be more responsive with a compliant force/torque sensor, but positional accuracy is reduced with compliance. An instrumented compliant wrist device can be used to achieve both responsive force control and accurate position control.

The wrist is connect in series between the end of the robot and the tool, and is designed to partially surround the tool, thus reducing the distance between the end of the robot and the end of the tool. The wrist device uses rubber elements for compliance and damping, and a serial linkage, with potentiometers at each joint, is used for sensing the deflections produced in the wrist.

This document describes the newest version of the instrumented compliant wrist, including modifications and improvements to the wrist described in "Design of a Tool Surrounding compliant Instrumented Wrist", available as tech report MS-CIS-91-30, GRASP LAB 258 from the University of Pennsylvania. Changes include a more protective sensing linkage structure and improved electronics. The compliance, kinematcis, and accuracy of the wrist are presented. Also, software for determining the wrist transform, and plans for the wrist are given.

A Proposal Concerning The Analysis Of Shadows In Images By An Active Observer (Dissertation Proposal)

Gareth D. Funka-Lea

MS-CIS-92-78

GRASP LAB 335

Shadows occur frequently in indoor scenes and outdoors on sunny days. Despite the information inherent in shadows about a scene's geometry and lighting conditions, relatively little work in image understanding has addressed the important problem of recognizing shadows. This is an even more serious failing when one considers the problems shadows pose for many visual techniques such as object recognition and shape from shading. Shadows are difficult to identify because they cannot be infallibly recognized until a scene's geometry and lighting are known. However, there are a number of cues which together strongly suggest the identification of a shadow. We present a list of these cues and methods which can be used by an *active* observer to detect shadows. By an active observer, we mean an observer that is not only mobile, but can extend a probe into its environment. The proposed approach should allow the extraction of shadows in real time. Furthermore, the identification of a shadow should improve with observing time. In order to be able to identify shadows without or prior to obtaining information about the arrangement of objects or information about the spectral properties of materials in the scene, we provide the observer with a probe with which to cast its own shadows. Any visible shadows cast by the probe can be easily identified because they will be new to the scene. These actively obtained shadows allow the observer to experimentally determine the number and location of light sources in the scene, to locate the cast shadows, and to gain information about the likely spectral changes due to shadows. We present a novel method for locating a light source and the surface on which a shadow is cast. It takes into account errors in imaging and image processing and, furthermore, it

takes special advantage of the benefits of an active observer. The information gained from the probe is of particular importance in effectively using the various shadow cues. In the course of identifying shadows, we also present a new modification on an image segmentation algorithm. Our modification provides a general description of color images in terms of regions that is particularly amenable to the analysis of shadows.

Striaght Line Walking Animation Based On Kinematic Generalization That Preserves The Original Characteristics

Kyeongseok Ko
Norman I. Badler
MS-CIS-92-79
GRAPHICS LAB 50

The most prominent problems in utilizing the rotoscopy data for human walking animation can be summarized into two: Preservation of the original motion characteristics in the *generalization* process and the Constraint Satisfaction.

Generalization is the process of producing the step of an arbitrary body and step length out of the original measured step which is of one particular subject and step length. If we lose much of the original style in the generalization, it would be meaningless to use the measured data. We present a generalization technique that keeps the original motion characteristics as much as possible.

Two types of generalization are considered. The one is the *body condition generalization*, which handles the differences between the two bodies. The ratio between the corresponding segments of the two bodies may not be uniform, which makes this generalization complicated. The other one is the *step length generalization*, which provides the steps with different step lengths of the same subject. These two generalizations are combined together to generate a step of arbitrary subject and step length.

The constraint satisfaction is enforced inside of our generalization process. Therefore the only thing that concerns us is the *quality* of the generalization. In our work, the preservation of the original characteristics is considered as the criteria determining the quality of the generalization. We prove that our generalization scheme actually preserves the characteristics of the original walk.

Proceedings of the Workshop on Linear Logic and Logic Programming (Washington, DC)

Edited by Dale Miller
MS-CIS-92-80
LINC LAB 238

Declarative programming languages often fail to effectively address many aspects of control and resource management. Linear logic provides a framework for increasing the strength of declarative programming languages to embrace these aspects. Linear logic has been used to provide new analyses of Prolog's operational semantics, including left-to-right/depth-first search and negation-as-failure. It has also been used to design new logic programming languages for handling concurrency and for viewing program clauses as (possibly) limited resources. Such logic programming languages have proved useful in areas such as databases, object-oriented programming, theorem proving, and natural language parsing.

This workshop is intended to bring together researchers involved in all aspects of relating linear logic and logic programming. The proceedings includes two high-level overviews of linear logic, and six contributed papers.

Workshop organizers: Jean-Yves Girard (CNRS and University of Paris VII), Dale Miller (chair, University of Pennsylvania, Philadelphia), and Remo Pareschi, (ECRC, Munich).

More On Goal-Directed Diagnosis

Ron Rymon
MS-CIS-92-81
LINC LAB 239

In many diagnosis-and-repair domains, diagnostic reasoning cannot be abstracted from repair actions, nor from actions necessary to obtain diagnostic information. We call these exploratory-corrective domains. In TraumAID 2.0, a consultation system for multiple trauma management, we have developed and implemented a framework for reasoning in such domains which integrates diagnostic reasoning with planning and action. In this paper, we present Goal-Directed Diagnosis (GDD), the diagnostic reasoning component of this framework. Taking the view that a diagnosis is only worthwhile to the extent that it can affect subsequent decisions, GDD focuses on the formation of appropriate goals for its complementary planner.

Active Color Image Analysis For Recognizing Shadows

Gareth Funka-Lea
Ruzena Bajcsy
MS-CIS-92-82
GRASP LAB 336

Many existing computer vision modules assume that shadows in an image have been accounted for prior to their application. In spite of this, relatively little work has been done on recognizing shadows or on recognizing a single surface

material when directly lit and in shadow. This is in part because shadows cannot be infallible recognized until a scene's lighting and geometry are known. However, color is a strong cue to the presence of shadows. We present a general color image segmentation algorithm whose output is amenable to the recovery of shadows as determined by an analysis of the physics of shadow radiance. Then, we show how an observer that can cast its own shadows can infer enough information about a scene's illumination to refine the segmentation results to determine where the shadows in the scene are with reasonable confidence. Having an observer that can actively cast shadows frees us from restrictive assumptions about the scene illumination or the reliance on high level scene knowledge. We present results of our methods on images of complex indoor and outdoor scenes.

Learning for Coordination of Vision and Action
Marcos Salganicoff
(University of Pennsylvania)
Ruzena Bajcsy
(University of Pennsylvania)
Tom Mitchell
(Carnegie Mellon University)
MS-CIS-92-84
GRASP LAB 337

We define the problem of visuomotor coordination and identify bottleneck problems in the implementation of general purpose vision and action systems. We conjecture that machine learning methods provide a general purpose mechanism for combining specific visual and action modules in a task-independent way. We also maintain that successful learning systems reflect realities of the environment, exploit context information, and identify limitations in perceptua l algorithms which cannot be captured by the designer. We then propose a multi-step find-and-fetch mobile robot search and retrieval task. This task illustrates where current learning approaches provide solutions and where future research opportunities exist.

Proceedings of the Workshop on the λProlog Programming Language
(University of Pennsylvania)
Edited by Dale Miller
MS-CIS-92-86

Learning and Forgetting for Perception-Action: A Projection Pursuit and Density Adaptive Approach (Dissertation)
Marcos Salganicoff
MS-CIS-92-87
GRASP LAB 338

We study learning of perception-action relations using visually-driven grasping as an example task. The well-established technique of non-parametric Projection Pursuit Regression (PPR) is used to accomplish reinforcement learning by searching for projections of high-dimensional data sets that capture invariants in the distribution of reinforcement in the parameter-space. The variable resolution 2^k-tree, a generalized quadtree, is used to represent perception-action maps based on the resulting reinforcement regression function.

We also pursue the following problem: how can we use human expertise and insight into grasping to train a system to select gripper approach directions and orientations for grasping, and then have it verify and adapt its skills through trial and error? To accomplish this learning we develop a new *Density Adaptive* Reinforcement Learning algorithm. This algorithm uses statistical tests to identify regions of the attribute space in which the dynamics of the task change and the density of exemplars is high. It concentrates the building of high-resolution descriptions in those areas.

In order to adapt the default rules to those necessary for the robot, it is necessary for the system to be able to forget previous experiences that no longer reflect the behavior of the world. A general purpose *Density Adaptive* forgetting algorithm has been developed that can be used as a front-end for a variety of learning methods. Additionally, by setting the forgetting parameters appropriately, an upper bound on the number of exemplars stored in the system may also be selected. This is important since all memory-based learning systems have finite memory in practice.

The approach is verified through simulation and experimentation. A robotic system incorporating two robots with a gripper, compliant instrumented wrist, arm, camera and laser scanner is used for experimentation. Since trial and error learning processes imply that failures will occur, the mechanics of the untrained robotic system must be able to tolerate mistakes during learning and not be damaged by excessive forces. We address this by the use of an instrumented, compliant robot wrist that controls impact forces.

Semantic Representations and Query Languages for Or-sets

Leonid Libkin
Limsoon Wong
MS-CIS-92-88
LOGIC & COMPUTATION 53

Or-sets were introduced by Imielinski, Naqvi and Vadaparty for dealing with liimited forms of disjunctive information in database queries. Independently, Rounds used a similar notion for representing disjunctive and conjunctive information in the context of situation theory. In this paper we formulate a query language with adequate expressive power for or-sets. Using the notion of normalization of or-sets, queries at the "structural" and "conceptual" levels are distinguished. Losslessness of normalization is established for a large class of queries. We have obtained upper bounds for the cost of normalization. An approach related to that of rounds is used to provide semantics for or-sets.

Syntactic Locality and Tree Adjoining Grammar: Grammatical, Acquisition and Processing Perspectives (Dissertation)

Robert Evan Frank
MS-CIS-92-89
LINC LAB 241

It has been widely recognized that the relations huuman grammar exploits are sensitive to constraints on structural locality. Indeed, much research in generative syntax has forcused on the precise characterization of the locality conditions that grammatical processes respect. In this dissertation, I propose that locality reflects the underlying formal system with which grammars are represented. In particular, I argue that the formalism of Tree Adjoining Grammar (TAG) is the appropriate meta-language for grammatical principles. TAG privdes a mechanism for composing phrase structure representations from small structural domains, and in so doing, restricts the class of possible grammatical principles to those expressible over these domains. Under this view, the existence of locality conditions is not directly stipulated, but instead follows from the representational machinery which the formal grammar makes available. I consider evidence from three domains of linguistic inquiry which provide convergent support for this view.

I first address the problem of constructing a grammatical theory in the principles and parameters fromaework in the context of the TAG formalism. I develop a substantive theory of the atomic objects of TAG, *elementary trees*, and argue for a Condition on Elementary Tree Minimality (CETM) which restricts the domain of an elementary tree. the CETM combined with TAG versions of the Projection Principle and the Empty Category Principle yields elegant analyses of a range of contructions including raising, copular sentences, wh-movement and gerunds.

Next, I turn to the domain of parsing and demonstrate how a TAG-based theory of grammar resolves the apparently incompatible demands of grammatical transparency and computational efficiency. The model I develop operates incrementally by processing in elementary tree-sized chunk,s as determined by the CETM. This view of incrementality is supported by psycholinguistic evidence concerning the closure points of syntactic processing.

Finally, in the domain of acquisitions, I argue that the assumption that children cannot perform the more complex operations of TAG, adjoining, explains relative difficulties in the acquisition of seemingly disparate constructions including wh-movement, control, raising, and relative clauses. This represents a new kind of explanation for the time course of acquisition through properties of formal and computational complexity.

Physics-Based Modeling Of Nonrigid Objects For Vision and Graphics (Dissertation)

Dimitri N. Metaxas
MS-CIS-92-90
GRAPHICS LAB 51

This thesis develops a physics-based framework for 3D shape and nonrigid motion modeling for computer vision and computer graphics. In computer vision it addresses the problems of complex 3D shape representation, shape reconstruction, quantitative model extraction from biomedical data for analysis and visualization, shape estimation, and motion tracking. In computer graphics it demonstrates the generative power of our framework to synthesize constrained shapes, nonrigid object motions and object interactions for the purposes of computer animation.

Our framework is based on the use of a new class of dynamically deformable primitives which allow the combination of global and local deformations. It incorporates physical constraints to compose articulated models from deformable primitives and provides force-based techniques for fitting such models to sparse, noise-corrupted 2D and 3D visual data. The framework leads to shape and nonrigid motion estimators that exploit dynamically deformable models to track moving 3D objects from time-varying observations.

We develop models with global deformation parameters which represent the salient shape features of natural parts, and local deformation parameters which capture shape details. In the context of computer graphics, these models represent the physics-based marriage of the parameterized and free-form modeling paradigms. An important benefit of their

global/local descriptive power in the context of computer vision is that it can potentially satisfy the often conflicting requirements of shape reconstruction and shape recognition.

The Lagrange equations of motion that govern our models, augmented by constraints, make them responsive to externally applied forces derived from input data or applied by the user. This system of differential equations is discretized using finite element methods and simulated through time using standard numerical techniques. We employ these equations to formulate a shape and nonrigid motion estimator. The estimator is a continuous extended Kalman filter that recursively transforms the discrepancy between the sensory data and the estimated model state into generalized forces. These adjust the translational, rotational, and deformational degrees of freedom such that the model evolves in a consistent fashion with the noisy data.

We demonstrate the interactive time performance of our techniques in a series of experiments in computer vision, graphics, and visualization.

Negotiation, Feedback, and Perspective Within Natural Language Generation (Dissertation)

Robert Rubinoff
MS-CIS-92-91
LINC LAB 242

This thesis is an investigation of how natrual language generation can take advantage of the ways that language use goes beyond simple, straightforward transmission of information. The two main constributions of the work are the use of annotations to relate linguistic choice to text planning and the use of perspective to model dependence and effects on attitude and context. Placing annotations on linguistic options allow the text planner to detect and respond to interactions between linguistic choices and the communicative plan driving those choices, without conflating the planning and linguistic levels of decision-making. The explicit modeling of perspective and perspective shifts allws the perspective to influence the generator's choices; in particular, the generator can take into account the ways that particular linguistic choices can affect or alter the subsequent perspective. The use of these techniques has been demonstrated in the IGEN implementation, resulting in a generator that can take advantage of the flexibility of natural language to frame its output in a way that reinforces the underlying goals driving the generation.

Robust Signal Restoration and Local Estimation of Image Structure

Visa Koivunen
MS-CIS-92-92
GRASP LAB 339

A class of nonlinear regression filters based on robust theory is introduced. the goal of the filtering is to restore the shape and preserve the details of the original noise-free signal, while effectively attenuating both impulsive and nonimpulsive noise. The proposed filters are based on robust Least Trimmed Squares estimation, where very deviating samples do not contribute to the final output. Furthermore, if there is more than one statistical population present in the processing windown the filter is very likely to select adaptively the samples that represent the majority and uses them for computing the output. We apply the regression filters on geometric signal shapes which can be found, for example, in range images. The proposed methods are also useful for extracting the trend of the signal without loosing important amplitutde information. We show experimental results on restroation of the original signal shape using real and synthetic data and both impulsive and nonimpulsive noise.

In addition, we apply the robust approach for describing local image structure. We sue the method for estimating spatial properties of the image in a local neighborhood. Such properties can be used for example, as auniformity predicate in the segmentation phase of an image understanding task. the emphasis is on producing reliable results even if the assumptions on noise, data and model are not completely valid. The experimental results provide information about the validity of those assumptions. Image describption results are shown using synthetic and real data, various signal shapes and impulsive and nonimpulsive noise.

Robust Location Estimation for MLR and Non-MLR Distributions (Dissertation)

Gerda L. Kamberova
MS-CIS-92-93 GRASP LAB 340

We study the problem of estimating an unknown parameter θ from an observation of random variable $Z = \theta + V$. this is the location data model' V is random noise with absolutely continuous distribution F, independent of θ. The distribution F belongs to a given uncertainty class of distributions $\mathcal{F}, |\mathcal{F}| \geq 1$. We seek robust minimax decision rules for estimating the location parameter θ. The parameter space is restricted - a known compact interval. The minimax risk is evaluated with respect to a zero-one loss function with a given error-tolerance e. the zero-one loss uniformly penalizes estimates which differ from the true parameter by more than the threshold e (these are unacceptable errors). the minimax criterion with zero-one loss is suitable for modeling problems for which it is desirable to minimize the maximum probability of getting unacceptable errors. As a consequence of this approach we obtain fixed size confidence intervals with highest probability of coverage.

We consider the distribution-dependent function $\frac{(x+2e)}{f(x)}$, where e is the error-tolerance and f is the noise density. We distinguish two different types of problems (involving two different types of distributions) based on behavior of this ratio: (I). Type $\mathcal{M}$-problems ($\mathcal{M}$-distributions) are characterized by a strictly monotone decreasing ratio; the minimax rules for $\mathcal{M}$-problems are admissible. They are monotone nondecreasing with a very simple structure – continuous, piecewise-linear. the class of $\mathcal{M}$-problems includes, but is not limited to, the distributions with monotone likelihood ratio (MLR) and non-MLR mixtures of normal distributions. (II). Type $\mathcal{NM}$-problems ($\mathcal{NM}$-distributions) are characterized by nonmonotone ratios; the minimax rules for these problems are in general nonmonotone.

The problem domain of low-level sensor fusion provides the motivation for our research. We examine sensor fusion problems for location data models using statistical decision theory. the decision-theoretic results wee obtain are sued for: (i) a robust test of the hypothesis that data from different sensors are consistent: and (ii) a robust procedure for combining the data which pass this preliminary consistency test.

A Gathering and Shooting Progressive Refinement Radiosity Method

Min-Zhi Shao
Norman I. Badler
MS-CIS-93-03
GRAPHICS LAB 52

This paper presents a gathering and shooting progressive refinement radiosity method. Our method integrates the iterative process of light energy gathering used in the standard full matrix method and the iterative process of light energy shooting used in the conventional progressive refinement method. As usual, in each iteration, the algorithm first selects the patch which holds the maximum unprocessed light energy in the environment as the shooting patch. But before the shooting process is activated, a light energy gathering process takes place. In this gathering process, the amount of the unprocessed light energy which is supposed to be shot to the current shooting patch from the rest of the environment in later iterations is pre-accumulated. In general, this extra amount of gathered light energy is far from trivial since it comes from every patch in the environment from which the current shooting patch can be seen. However, with the reciprocity relationship for form-factors, still only one hemi-cube of the form-factors is needed in each iteration step. Based on a concise record of the history of the unprocessed light energy distribution in the environment, a new progressive refinement algorithm with revised gathering and shooting procedures is then proposed. With little additional computation and memory usage compared to the conventional progressive refinement radiosity method, a solid convergence speedup is achieved. This gathering and shooting approach extends the capability of the radiosity method in accurate and efficient simulation of the global illuminations of complex environments.

An Algorithm for VP Ellipsis

Daniel Hardt
MS-CIS-93-050
LINC 244

An algorithm is proposed to determine antecedents for VP ellipsis. The algorithm eliminates impossible antecedents, and then imposes a preference ordering on possible antecedents. The algorithm performs with 94% accuracy on a set of 304 examples of VP ellipsis collected from the Brown Corpus. The problem of determining antecedents for VP ellipsis has received little attention in the literature, and it is shown that the current proposal is a significant improvement over alternative approaches.

VP Ellipsis and Contextual Interpretation

Daniel Hardt
MS-CIS-93-06
LINC LAB 245

A computational account of VP ellipsis is described, in which VP's are represented in the discourse model as contextually dependent semantic objects. It is argued that this approach can handle examples that are not allowed by alternative accounts. An implementation is defined in terms of extensions to the Incremental Interpretation System. The treatment of VP ellipsis is analogous to that of pronominal anaphora. It is suggested that the recency and salience constraints commonly thought to apply to pronominal anaphora might apply in a similar way to VP ellipsis.

A State Minimization Algorithm for Communicating State Machines with Arbitrary Data Space

Inhye Kang
Insup Lee
MS-CIS-93-07
LOGIC & COMPUTATION 55
DISTRIBUTED SYSTEMS LAB 13

A fundamental issue in the automated analysis of communicating systems is the efficient generation of the reachable state space. Since it is not possible to generate all the reachable states of a system with an infinite number of states, we need a way to combine sets of states. In this paper, we describe communicating state machines with data variables, which we use to specify concurrent systems. We then present an algorithm that constructs the minimal reachability graph of a labeled transition system with infinite data values. Our algorithm clusters a set of states that are bisimilar into an equivalent class. We include an example to illustrate our algorithm and identify a set of sufficient conditions that guarantees the termination of the algorithm.

A Process Algebraic Approach To The Specification and Analysis of Resource-Bound Real-time Systems

Insup Lee
(University of Pennsylvania)
Patrice Brémond-Grégoire
(University of Pennsylvania)
Richard Gerber
(University of Maryland)
MS-CIS-93-08
LOGIC & COMPUTATION 56
DISTRIBUTED SYSTEMS LAB 14

There has recently been significant progress in the development of timed process algebras for the specification

and analysis of real-time systems. This paper describes a timed process algebra called ACSR. ACSR supports synchronous timed actions and asynchronous instantaneous events. Timed actions are used to represent the usage of resources and to model the passage of time. Events are used to capture synchronization between processes. To be able to accurately specify real systems, ACSR supports a notion of priority that can be used to arbitrate among timed actions competing for the use of resources and among events that are ready for synchronization. The paper also includes a brief overview of other timed process algebras and discusses similarities and differences between them and ACSR.

Fast Parallel Deterministic and Randomized Algorithms for Model Checking

Insup Lee
Sanguthevar Rajasekaran
MS-CIS-93-09
GRASP LAB 342

Model checking is a powerful technique for verification of concurrent systems. One of the potential problems with this technique is state space explosion. There are two ways in which one could cope with state explosion: reducing the search space and searching less space. Most of the existing algorithms are based on the first approach.

One of the sccuessful approach for reducing search space uses Binary Decision Diagrames (BDDs) to represent the system. Systems with a large number of states (of the order of 5 x 10^{20}) have been thus verified. But there are limitations to this heuristic approach. Even systems of reasonable complexity have many more states. Also, the BDD apprach might fail even on some simple systems. In this paper we propose the use of parallelism to extend the applicability of BDDs in model checking. In particular we present very fast algorithms for model checking the employ BDDs. The algorithms presented are much faster than the best known previous algorithms. We also describe searching less space as an attractive approach to model checking. In this paper we demonstrate the power of this approach. We also suggest the use of randomization in the design of model checking algorithms.

Selection, Routing and Sorting On The Star Graph

Sanguthevar Rajasekaran
David S.L. Wei
MS-CIS-93-10
GRASP LAB 343

We consider the problems of selection, routing and sorting on an n-star graph (with $n!$ nodes), an interconnection network which has been proven to possess many special properties. We identify a tree like subgraphc (which we call as a '$(k, 1, k)$ chain network') of the star graph which enables us to design efficient algorithms for the above mentioned problems.

We present an algorithm that performs a sequence of n prefix computations in $O(n^2)$ time. This algorithm is used as a subroutine in our other algorithms. In additonal we offer an efficient deterministic sorting algorithm that runs in $O(n^3 \log n)/2$ steps. Though an algorithm with the same time bound has been proposed before, our algorithm is very simeple and is based on a different approach. We also show that sorting can be performed on the $n$star graph in time $O(n^3)$ and that selection of a set of uniformaly distributed n keys can be performed in $O(n^2)$ time with high probability. Finally, we also present a deterministic (non oblivious) routing algorithm that realizes any permutation in $O(n^3)$ steps on the n-star graph.

There exists an algorithm in the literature that can perform a single prefix computation in $O(n \log n)$ time. The best known previous algorithm for sorting has a run time of $O(n^3 \lg n)$ and is deterministic. To our knowledge, the problem of selection has not been considered before on the star graph.

Gigabit Telerobitcis: Applying Advanced Information Infrastructure

Ruzena Bajcsy
David J. Farber
Richard P. Paul
Jonathan M. Smith
MS-CIS-93-11
GRASP LAB 344
DISTRIBUTED SYSTEMS LAB 15

Advanced manufacturing concepts such as "Virtual Factories" use an information infrastructure to tie together changing groups of specialized facilites to agile manufacturing systems. A necessary element of such systems is the ability to *teleoperate* machines, for example telerobotic systems with full-capability sensory feeback loops. We have identified three network advances needed for splitting robotic *control* from robotic *function*: increased bandwidth, decreased error rates, and support for isochronous traffic. These features are available in the Gigabit networks under the development at Penn and elsewhere.

A number of key research questions are posed by gigabity telerobotics. Tehre are issues in network topology, robot control and distributed system software, packaging and transport of sensory data (including wide-area transport), and performance implications of architectural choices using measures such as cost, response time, and network utilization.

We propose the explore these question *experimentally* in

a joint research effort combining the Distributed Systems Laboratory and the General Robotics and Sensory Perception (GRASP) Laboratory at the University of Pennsylvania. The proposed experiments should provide important early results. A detailed research program is described.

A Kinematic Generalization of Rotoscoped Human Walking

Hyeongseok Ko
Norman I. Badler
MS-CIS-93-12
GRAPHICS LAB 54

The most prominent problems in utilizing the rotoscopy data for human walking animation can be summarized as preservation of the original motion characteristics in the *generalization* process and *constraint satisfaction*. Generalization is the process of producing the step of an arbitrary body and step length from the original measured step of one particular subject and step length. If we lose much of the original style in the generalization, it would be meaningless to use the measured data. We present a generalization technique that keeps the original spatial motion characteristics as much as possible. Two types of generalization are considered. One is the *anthropometry generalization*, which handles the non-uniform segment length ratio differences between the two bodies. The other is the *step length generalization*, which changes the steps to different step lengths of the same subject. These two generalizations are combined together to generate a step of arbitrary subject and step length. The constraint satisfaction is enforced within the generalization process.

Curved Path Human Locomotion that Handles Anthropometrical Variety

Hyeongseok Ko
Norman I. Badler
MS-CIS-93-13
GRAPHICS LAB 53

Human locomotion simulation along a curved path is presented. The process adds a small constant cost ($O(1)$) to any pre-existing straight line walking algorithm. The input curve is processed by the *foot print generator* to produce a foot print sequence. The resulting sequence is scanned by the *walking motion generator* that actually generates the poses of the walking that realizes such foot prints. The two primitives **INITIALISE_STEP** and **ADVANCE_STEP** are used for walking motion generation. **INITIALIZE_STEP** is activated with the input parameters *walker*, *next_foot_print*, *left_or_right*, and *step_duration*, just before each step to precompute the trajectories of the center of the body and the ankles. **ADVANCE_STEP** is called with a *normalized time* to generate the actual pose at that moment. The normalized time is a logical time, covering zero to one during a complete step.

Operator Interaction and Teleprogramming for Subsea Manipulation

Craig Sayers
Richard Paul
Max Mintz
MS-CIS-93-15
GRASP LAB 344

The teleprogramming paradigm has been proposed as a means to efficiently perform teleoperation in the subsea environment via an acoustical link. In such a system the effects of both limited bandwidth channels and delayed communications are overcome by transmitting not Cartesian or joint level information but rather symbolic, error-tolerant, program instructions to the remote site. The operator interacts with a virtual reality of the remote site which provides immediate visual and kinesthetic feedback. The uncertainty in this model can be reduced based on information received from the slave manipulator's tactile contact with the environment. It is suggested that the current state of the model be made available to the operator via a graphical display which shows not only the position of objects at the remote site but also, through the use of color clues, the uncertainty associated with those positions. The provision of uncertainty information is important since it allows the operator to compromise between speed and accuracy. An additional operator aid, which we term synthetic fixturing, is proposed. Synthetic fixtures provide the operator of the teleprogramming system with the teleoperation equivalent of the "snap" commands common in computer aided design programs. By guiding the position and/or orientation of the master manipulator toward specific points, lines or planes the system is able to increase both the speed and precision with which the operator can control the slave arm without requiring sophisticated hardware.

DNA Workbench

James Tisdall
MS-CIS-93-38
LOGIC & COMPUTATION 61

DNA WorkBench is a program for working with DNA, RNA, and protein sequences. It is designed to solve several problems that arise in two domains. The first domain is that of the algorithm designer and implementor who is working in the emerging field of computational biology. The second domain is that of the worker in a genetics laboratory, who needs frequently to turn to the computer to perform

analysis on existing or newly acquired nucleotide or protein sequences. The problems encountered in these two domains overlap to a considerable extent. The problems, and how they are addressed by **DNA WorkBench**, are discussed within.

DNA WorkBench addresses both of these groups with one program. In this way, new problems that require new algorithms can be quickly brought from a theoretical solution to an implementation and to the laaboratory workbench. This rapid transfer from research to development to the field is essential in a fast moving area such as biotechnology, by which term I mean to encompass such specialties as molecular biology, genetics, human gene therapy, and the current large-scale international sequencing and mapping of human DNA which has been organized as the Human Genome Project in the United States.

Generating Contextually Appropriate Intonation

Scott Prevost
Mark Steedman
MS-CIS-93-39
LINC LAB 247

One source of unnaturalness in the output of text-to-speech systems stems from the involvement of algorithmically generated default intonation contours, applied under minimal control from syntax and semantics. It is a tribute both to the resilience of human language understanding and to the ingenuity of the inventors of these algorithms that the results are as intelligible as they are. However, the result is very frequently unnatural, and may on occasion mislead the hearer. This paper extends earlier work on the relation between syntax and intonation in language understanding in Combinatory Categorial Grammar (CCG). A generator with a simple and domain-independent discourse model can be used to direct synthesis of intonation contours for responses to data-base queries, to convey distinctions of contrast and emphasis determined by the discourse model.

Model Based Teleoperation To Eliminate Feedback Delay
NSF Grant BCS89-01352 - 3rd Report

Richard P. Paul
Thomas Lindsay
Craig Sayers
Matthew Stein
Desikachar Venkatesh
MS-CIS-93-40
GRASP LAB 345

We are conducting research in the area of teleoperation with feedback delay. Significant delays occur when performaing space teleoperation from the earth as well as in subsea teleoperation where the operator is typically on a surface vellel and communication is via acoustic links. These delays make teleoperation extremely difficult and lead to very low operator productivity.

We have combined computer graphics with manipulator programming to provide a solution to the delay problem. A teleoperator master arm is interfaced to a graphical simulation of the remote environment. Synthetic fixtures are used to guid the operators motions and to provide kinesthetic feedback. The operator's actions are monitored and used to generate symbolic motion commands for transmission to, and execution by, the remote slave robot. While much of a task proceeds error free, when an error does occur, the slave system transmits data back to the master environment where the opertor can then experience the motion of the slave manipulator in actural task execution. We have also provided for the use of tools such as an impact wrench and a winch at the slave site. In all cases the tools are unencumbered by sensors: the slave uses a compliant instrumented wrist to monitor tool operation in terms of resulting motions and reaction forces.

MS-CIS-93-41 Val's Paper LOGIC & COMPUTATION 62

An SE-tree based Characterization of the Induction Problem

Ron Rymon
MS-CIS-93-42
LINC LAB 248

Many induction programs use decision trees both as the basis for their search, and as a representation of their classifier solution. In this paper we propose a new structure, called SE-tree, as a more general alternative.

Physics-Based Modeling, Analysis and Animation

Ioannis A. Kakadiaris
MS-CIS-93-45
GRASP LAB 346

The idea of using physics-based models has received considerable interest in computer graphics and computer vision research the last ten years. The interest arises from the fact that simple geometric primitives cannot accurately represent natural objects. In computer graphics physics-based models are used to generate and visualize constrained shapes, motions of rigid and nonrigid objects and object interactions with the environment for the pur- poses of animation. On the other hand, in computer vision, the method applies to complex 3-D shape representation, shape recon- struction and motion estimation.

In this paper we review two models that have been used in computer graphics and two models that apply to both areas. In the area of computer graphics, Mill- er uses a mass-spring model to animate three forms of locomo- tion of snakes and worms. To overcome the problem of the multitude of degrees of freedom associated with the mass-spring lat- tices, Witkin and Welch present a geometric method to model global deformations. To achieve the same result Pentland and Horowitz in delineate the object motion into rigid and nonri- gid deformation modes. To overcome problems of these two last approaches, Metaxas and Terzopoulos in successfully combine local deformations with global ones.

Modeling based on physical principles is a potent technique for computer graphics and computer vision. It is a rich and fruitful area for research in terms of both theory and applications. It is important, though, to develop concepts, methodologies, and techniques which will be widely applicable to many types of applications.

An Active Approach to Characterization and Recognition of Functionality and Functional Properties

Luca Bogoni
Ruzena Bajcsy
MS-CIS-93-50
GRASP LAB 347

Functionality in an object can be defined as its applicability toward the accomplishment of a task. We emphasize and develop an interactive and performatory approach to functionality recovery from sensor data in the context of robotic manipulatory tasks. By analyzing interaction of tool and target object and manipulation tasks as goal–oriented recognition processes we propose to identify and characterize functionalities of objects. This interaction is not only a means of verification of the hypothesized presence of functionality in objects but also a way to actively and purposively recognize the object. The representation of functionality allows us to extend the recovery process to a hierarchy of functionalities allowing complex ones to be composed from simpler ones.

A formal model, based on Discrete Event Dynamic System Theory (DEDS), is introduced to define an interactive task for recovering and describing functionality. To observe and control the recovery process we introduce the notion of piecewise observability of a task by different sensors. This allows the description of a dynamic system in which not all events nor the time of their occurrence may be predicted in advance. An experimental system, with both vision and force sensors, for carrying out the interactive functional recognition is described.

Elastically Deforming A Three-Dimensional Atlas To Match Anatomical Brain Images

Jim C. Gee
Martin Reivich
Ruzena Bajcsy
MS-CIS-93-53
GRASP LAB 348

To evaluate our system for elastically deforming a three-dimensional atlas to match anatomical brain images, six deformed versions of an atlas were generated. The deformed atlases were created by elastically mapping an anatomical brain atlas onto different MRI brain image volumes. The mapping matches the edges of the ventricles and the surface of the brain; the resultant deformations are propagated through the atlas volume, deforming the remainder of the structures in the process. The atlas was then elastically matched to its deformed versions. The accuracy of the resultant matches was evaluated by determining the correspondence of 32 cortical and subcortical structures. The system on average matched the centroid of a structure to within 1 mm of its true position and fit a structure to within 11% of its true volume. The overlap between the matched and true structures, defined by the ratio between the volume of their intersection and the volume of their union, averaged 66%. When the gray-white interface was included for matching, the mean overlap improved to 78%; each structure was matched to within 0.6 mm of its true position and fit to within 6% of its true volume. Preliminary studies were also made to determine the effect of the compliance of the atlas on the resultant match.

Modeling Articulated Figure Motion With Physically-and Physiologically-Based Constraints (Dissertation)

Philip L.Y. Lee
MS-CIS-93-54
GRAPHICS LAB 55

A methodology and algorithm are presented that generate motions imitating the way humans complete a task under various loading conditions. The path taken depends on "natural" parameters: the figure geometry, the given load, the final destination, and, especially, the *strength model* of the agent. Additional user controllable parameters of the motion are the *comfort* of the action and the *perceived exertion* of the agent. The algorithm uses this information to incrementally compute a motion path of the end-effector moving the load. It is therefore instantaneously adaptable to changing force, loading, and strength conditions. Various strategies are used to model human behavior (such as available torque, reducing moment and pull back) that direct the trajectories.

The strength model dictates acceptable kinematic postures. The resulting algorithm also offers torque control without the tedious user expression of driving forces under a dynamics model. The algorithm runs in near-real-time and offers an agent-dependent toolkit for fast path prediction. Examples are presented for various lifting tasks, including one- and two-handed lifts, and raising the body from a seated posture.

Intermittent Non-Rhythmic Human Stepping and Locomotion

Hyeongseok Ko

MS-CIS-93-55
GRAPHICS LAB 56

When humans need to get from one location to another, there are many occasions where non-rhythmic stepping (NRS) is more desirable than normal walking. This can be observed in performing tasks in a constricted work space. For the purpose NRS is considered as a variation of curved path walking. Four types of local adjustment are dealt with: forward, backward, lateral stepping, and turnaround. Combined with curved path walking, NRS provides a very useful tool for animating human locomotion behaviors. In the lower body motion, the trajectory of the hip, angular trajectory of the feet, and the trajectory of the swing ankle during the swing phase determine the basic outline of an NRS. These trajectories are precomputed at the start of each step. The stepping process is called with a *normalized time* to generate the actual pose of the NRS at that moment. the normalized time is a logical time, covering zero to one during a complete step.

Search Plans (Dissertation Proposal)

Michael Moore

MS-CIS-93-56
LINC LAB 250

People often do not know where things are and have to look for them. This thesis presents a formal model suitable for reasoning about how to find things and acting to find them, which I will call "search behavior". Since not knowing location of something can prevent an agent from reaching its desired goal, the ability to plan and conduct a search will be argued to increase the variety of situations in which an agent can succeed at its chosen task.

Searching for things is a natural problem that arises when the blocks world assumptions (which have been the problem setting for most planning research) are modified by providing the agent only *partial* knowledge of its environment. Since the agent does not know the total world state, actions may *appear* to have nondeterministic effects. The significant aspects of the search problem which differ from previously studied planning problems are the acquisition of information and iteration of similar actions while exploring a search space.

Since introduction of the situation calculus [?], various systems have been proposed for *representing* and *reasoning* about actions which involve knowledge acquisition and iteration, including Moore's work on the interaction between knowledge and action [?]. My concern with searching has to do with a sense that Moore's knowledge preconditions are overly restrictive. Morgenstern [?] examined ways to weaken knowledge preconditions for an individual agent by relying on the knowledge and abilities of other agents. Lesperance's research [?] on indexical knowledge is another way of weakening the knowledge preconditions. I am trying to reduce the *amount* of information an agent must know (provided they can search a known search space). If you dial the right combination to a safe it will open, whether or not you knew in advance that it *was* the right combination. Search is a way to guarantee you will eventually dial the right combination. So what I am exploring is how to systematically construct a search that will use available knowledge to accomplish something the agent does not currently know enough to do directly. Such systems can be used to infer properties of plans which have already been constructed, but do not themselves *construct plans* for complex actions.

I claim it is possible for automated agents to engage in search behavior. Engaging in search behavior consists in recognizing the need for a search, constructing an effective plan, and then carrying out that plan. Expressing such a plan and reasoning about its effectiveness requires a representation language. I will select a representation language based on criteria derived from analyzing the search planning problem. Each of the three components of a system for engaging in search behavior will be designed and implemented to demonstrate that an automated agent can find things when it needs to.

Goal-Directed Diagnosis-Diagnostic Reasoning in Exploratory-Corrective Domains

Ron Rymon

MS-CIS-93-57
LINC LAB 251

In many diagnosis-and-repair domains, diagnostic reasoning cannot be abstracted from repair actions, nor from actions necessary to obtain diagnostic information. We call these exploratory-corrective domains. In TraumAID 2.0, a consultation system for multiple trauma management, we have developed and implemented a framework for reasoning in such domains which integrates diagnostic reasoning with

planning and action. In this paper, we present Goal-Directed Diagnosis (GDD), the diagnostic reasoning component of this framework. Taking the view that a diagnosis is only worthwhile to the extent that it can affect subsequent decisions, GDD focuses on the formation of appropriate goals for its complementary planner.

The Soundness and Completeness of ACSR (Algebra of Communicating Shared Resources)

Patrice Brémond-Grégoire
Jin-Young Choi
Insup Lee
MS-CIS-93-59
GRASP LAB 350

Recently, significant progress has been made in the development of timed process algebras for the specification and analysis of real-time systems; one of which is a timed process algebra called ACSR supports synchronous timed actions and asynchronous instantaneous events. Timed actions are used to represent the usage of resources and to model the passage of time. Events are used to capture synchronization between processes. The be able to specify real systems accurately, ACSR supports a notion of priority that can be used to arbitrate among timed actions competing for the use of resources and among events that are ready for synchronization. Equivalence between ACSR terms is defined in terms of strong bisimulation. The paper contains a set of algebraic laws that are proven sound and complete for finite ACSR agents. This paper contains the soundness and completeness proofs of the ACSR laws reported in an earlier report (MS-CIS-93-08)

Instructions, Intentions and Expectations

Bonnie Webber
Norman Badler
Barbara DiEugenio
Christopher Geib
Libby Levinson
Michael Moore
MS-CIS-93-61
LINC LAB 252

This is a short position paper on what we have learned about the relationship between language and behavior from an on-going attempt to enable people to use Natural Language instructions to tell animated human figures what to do.

We view instructions as texts intended to be understood in context, produced by an instructor who knows more than the agent. (The latter means that instructions are worth trusting, even if the world initially provides no corroborating evidence.) With this view underlying the architecture of our Animation from Natural Language (*AnimNL*) system, the paper discusses two main things we have learned about the complex relationship between language and behavior:

1. *Intentions* formed in response to instructions influence agents' behavior at every level of decision-making, from language understanding to motor control.
2. *Expectations* formed in response to instructions influence agents' behavior — what they do and what they look for — over and beyond their current perceptions. As such, expectations from instructions complement information from the world in guiding an agent's behavior.

Queries on Databases with User-Defined Functions

Dan Suciu
MS-CIS-93-62
LOGIC & COMPUTATION 70

The notion of a database query is generalized for databases with user-defined functions. Then, we can prove that the computable queries coincide with those expressible by an extension of the relational machine, with oracles ([4]). This implies that any complete query language, extended with user-defined function symbols in a "reasonable" way, is still complete. We give an example of a complete query language with user-defined functions, and discuss its connections with object inventions.

SASS v.2.1 Anthropometric Spreadsheet and Database for the IRIS

Francisco Azuola
Teo Kok Hoon
Sussana Wei
MS-CIS-93-63
GRAPHICS LAB 57

It describes the usage of SASS, a spreadsheet-like system which allows flexible interactive access to all anthropometric variables needed to size a computer-based human figure, described structurally by a PEABODY file.

Data that may be accessed is organized into the following "groups": segment dimension ("girth"), joint limits, center of mass, and strength, all of which work based on statistical population data. SASS creates generic computer-based human figures based on this data.

SASS also is an anthropometric database and interactive query system that works upon anthropometric data of real individuals. Scaled computer-based human figures created by SASS can be displayed directly, and interactively changed, within the *Jack* software.

Multiple Representation Approach To Geometric Model Construction from Range Data

Visa Koivunen
Jean-Marc Vezien
Ruzena Bajcsy

MS-CIS-93-66
GRASP LAB 352

A method is presented for constructing geometric design data from noisy 3-D sensor measurements of physical parts. In early processing phase, RLTS regression filters stemming from robust estimation theory are used for separating the desired part of the signal in contaminated sensor data from undesired part. Strategies for producing a complete 3-D data set from partial views are studied. Surface triangulation, NURBS, and superellipsoids are employed in model construction to be able to represent efficiently polygonal shapes, free form surfaces and standard primitive solids. Multiple representations are used because there is no single representation that would be most appropriate in all situations. The size of the required control point mesh for spline description is estimated using a surface characterization process. Surfaces of arbitrary topology are modeled using triangulation and trimmed NURBS. A user given tolerance value is driving refinement of the obtained surface model. The resulting model description is a procedural CAD model which can convey structural information in addition to low level geometric primitives. The model is translated to IGES standard product data exchange format to enable data sharing with other processes in concurrent engineering environment. Preliminary results on view registration and integration using simulated data are shown. Examples of model construction using both real and simulated data are also given.

Simplifying Tool Usage In Teleoperative Tasks

Thomas Lindsay
Richard P. Paul

MS-CIS-93-68
GRASP LAB 353

Modern robotic research has presented the opportunity for enhanced teleoperative systems. *Teleprogramming* has been developed for teleoperation in time-delayed environments, but can also lead to increased productivity in non-delayed teleoperation.

Powered tools are used to increase the abilities of the remote manipulator. However, tools add to the complexity of the system, both in terms of control and sensing. Teleprogramming can be used to simplify the operators interaction with the manipulator/tool system. Further, the adaptive sensing algorithm of the remote site system (using an instrumented compliant wrist for feedback) simplifies the sensory requirements of the system. Current remote-site implementation of a teleprogramming tool-usage strategy that simplifies tool use is described in this document.

the use of powered tools in teleoperation tasks is illustrated by two examples, in using an air-powered impact wrench, and the other using an electric winch. both of these tools are implemented at our remote site workcell, consisting of a Puma 560 robot working on the task of removing the top of a large box.

Verb Phrase Ellipsis: Form, Meaning and Processing (Dissertation)

Daniel Hardt

MS-CIS-93-69
LINC LAB 255

The central claim of this dissertation is that an elliptical VP is a proform. This claim has two primary consequences: first, the elliptical VP can have no internal syntactic structure. Second, the interpretation of VP ellipsis must be governed by the same general conditions governing other proforms, such as pronouns. The basic condition governing the interpretation of a proform is that it must be semantically identified with its antecedent. A computational model is described in which this identification is mediated by store and retrieve operations defined with respect to a discourse model. Because VP ellipsis is treated on a par with other proforms, the ambiguity arising from "sloppy identity" becomes epiphenomenal, resulting from the fact that the store and retrieve operations are freely ordered.

A primary argument for the proform theory of VP ellipsis concerns syntactic constraints on variables within the antecedent. I examine many different types of variables, including reflexives, reciprocals, negative polarity items, and wh-traces. In all these cases, syntactic constraints are not respected under ellipsis. This indicates that the relation governing VP ellipsis is semantic rather than syntactic. In further support of the proform theory, I show that there is a striking similarity in the antecedence possibilities for VP ellipsis and those for pronouns.

Two computer programs demonstrate the claims of this dissertation. One program implements the semantic copying required to resolve VP ellipsis, demonstrating the correct set of possible readings for the examples of interest. The second program selects the antecedent for a VP ellipsis occurrence. This program has been tested

on several hundred examples of VP ellipsis, automatically collected from corpora.

Algebraic Characterization of Edible Powerdomains
Leonid Libkin
MS-CIS-93-70
LOGIC & COMPUTATION 71

Powerdomains like mixes, sandwiches, snacks and scones are typically used to provide semantics of collections of descriptions of partial data. In particular, they were used to give semantics of databases with partial information. In this paper we argue that to be able to put these constructions into the context of a programming languages it is necessary to characterize them as free (ordered) algebras. Two characterizations – for mixes and snacks – are already known, and in the first part of the paper we give characterizations for scones and sandwiches and provide an alternative characterization of snacks. The algebras involved have binary and unary operations and relatively simple equational theories. We then define a new construction, which is in essence all others put together (hence called salad and give its algebraic characterization. It is also shown how all algebras considered in the paper are related in a natural way, that is, in a way that corresponds to embeddings of their powerdomains. We also discuss some semantic issues such as relationship between the orderings and the semantics and justification for choosing the orderings. Finally, we outline prospects for further research.

Parallel Algorithms for Relational Coarsest Partition Problems
Sanguthevar Rajasekaran
Insup Lee
MS-CIS-93-71
GRASP LAB 354

relational Coarsest Partition Problems (RCPPs) play a vital role in verifying concurrent systems. It is known that RCPPS are $\mathcal{P}$-complete and hence it may not be possible to design polylog time parallel algorithms for these problems.

In this paper, we present two efficient parallel algorithms for RCPP, in which its associated label transition system is assumed to have m transitions and n states. The first algorithm runs in $O(n^{1+\epsilon})$ time using $\frac{m}{n^{\epsilon}}$ CREW PRAAM processors, for any fixed $\epsilon < 1$. This algorithm is analogous and optimal with respect to the sequential algorithm of Kanellakis and Smolka. the second algorithm runs in $O(n \log n)$ time using $\frac{m}{n}$ log n CREW PRAM processor. this algorithm is analogous and nearly optimal with respect to the sequential algorithm of Paige and Tarjan.

Fast Parallel Routing and Computation On Interconnection Networks (Dissertation)
David Shyue Liang Wei
MS-CIS-93-72
GRASP LAB 355

both parallel processing and artificial intelligence play important roles in computer science. The application of parallel processing in artificial intelligence is one of the possible approaches to realize an efficient and intelligent computer. This is a two-part thesis. The first part of this thesis investigates routing problems, which are central to parallel processing, for a class of interconnection networks called *leveled networks*, while the second part of the thesis makes use of these results to develop efficient parallel algorithms for some communication intensive artificial intelligence problems. Specifically, we look at the parsing problem for natural language grammars which is a fundamental problem in artificial intelligence. Our work goes beyond theoretical study on routing and parallel algorithms in the we also develop implementations of the algorithms on an actual parallel machine, viz., the *Connection Machine* (CM). this allows us to verify experimentally the performance predicted by theoretical analysis.

To date, much of the work on routing has virtually centered on constant degree networks with *logarithmic* (or even larger) diameter. In order to achieve faster communication, we initiate the study of routing on some nonconstant degree networks in *sublogarithmic* diameter. We also give a universal randomized optimal routing algorithm for a large class of interconnection networks, viz., leveled networks. These leveled networks can be of logarithmic, or sublogarithmic diameter. Further, we present algorithms for emulating *PRAMs*, and ideal shared memory model, on leveled networks. The emulation is optimal.

We also study parallel parsing algorithms. In particular, we consider tree adjoining grammars (TAGs). We give a parallel parsing algorithm on a 5-dimensional systolic array, which achieves optimal speed-up with respect to the best known sequential one. We also present an efficient algorithm for general TAGs on the CM. This implemented algorithm parallelizes the parsing in terms of the grammar size unlike previous parallel parsing algorithms which parallelize the parsing in terms of the input sentence. The former is highly desirable for the natural language processing because of the huge grammar size.

A Lower Bound Result For The Common Element Problem and Its Implication For Reflexive Reasoning

Paul Dietz
Danny Krizanc
Sanguthevar Rajasekaran
Lokendra Shastri

MS-CIS-93-73
GRASP LAB 356

In this paper we prove a lower bound of Ω (n log n) for the common element problem on two sets of size n each. Two interesting consequences of this lower bound are also discussed. In particular, we show that linear space neural network models that admit *unbalanced* rules cannot draw all inferences in time independent of the knowledge base size. We also show that the *join* operation in data base applications needs Ω (log n) time given only n processors.

Moving Posture Reconstruction From Perspective Projections of Jointed Figure Motion (Dissertation)

Jianmin Zhao

MS-CIS-93-74
GRAPHICS LAB 58

Our goal is to reproduce a human figure's motion with a computer simulated human figure: Given a sequence of perspective projections of a set of feature joints of the moving figure, we tried to recover the original 3D postures through an accurate human figure model and the continuity requirement (temporal coherence) in the sequence. Our approach follows two clues: Given the human figure model, the responsible posture for a frame is constrained by the projections of all the feature joints and, in this limited set of postures we can choose one based on the postures in the previous frames and the temporal coherence, unless there occurs a critical condition, when the projection ray of a feature is perpendicular to the link of which the feature is the distal end. Owing to the fast inverse kinematics algorithm we developed to solve the spatial constraints, we were able to exploit the temporal coherence in projection sequences of frequencies as high as 100 Hz. We used finite state automata to detect critical conditions, and developed various strategies to overcome special difficulties around critical frames. Furthermore, we investigated the impact of errors in linear measurements of body parts on the reconstruction process. Based on mathematical analysis, we proposed some heuristics to discover and recover from the possible modeling errors. To test the theory, we implemented an experimental system. By imposing the temporal coherence constraint whenever possible, this system responds to the incoming images almost linearly: Since the error-prone critical conditions are detected and handled at the very early stage, the system is able to do away with endless recursive backtracking so that only one level of roll-back is needed to handle a limited number of critical conditions whose chances of occurrence are independent of the sampling rate. The system admits generic human motion. It has been tested on synthesized images from actual 3D human motions. Since we knew the original motion, we were able to evaluate results quantitatively. It turned out that the reconstructed motions agreed with the original ones not only in general but also in fine details.

Control of Mechanical Systems With Rolling Contacts: Applications To Robotics (Dissertation)

Nilanjan Sarkar

MS-CIS-93-75
GRASP LAB 357

The problems of modeling and control of mechanical dynamic systems subject to rolling contacts are investigated. There are two important theoretical contributions in this dissertation. First, contact kinematic relationships up to second order are developed for two rigid bodies in point contact. These equations relate gross rigid body motion to the changes in the positions of the points of contact. Second, a unified approach to the control of mechanical systems subject to both holonomic and nonholonomic constraints is proposed. The basic approach is to extend the state-space to include, in the addition to the generalized coordinates and velocities, contact coordinates which describe the displacements of the contact points and their derivatives. This redundant state-space formulation provides a convenient way to specify output equations to control contact motion. The control problem is formulated as an affine nonlinear problem and a differential-geometric, control-theoretic approach is used to decouple and linearize such systems. It is shown that such a system, even though not input-state linearizable, is input-output linearizable. Further, the zero dynamics of such a system is shown to be Lagrange stable.

The proposed methodology is applied to three different robotic systems: (a) wheeled mobile robots, (b) two arms manipulating an object with rolling contact between each arm and the object, and (c) a single robot arm maintaining controlled contact against a moving environment. In each case, a nonlinear controller is designed to achieve the desired performances. For mobile robots, a new control algorithm called dynamic path following is proposed and shown to be quite effective and robust. In the con-

text of two arm manipulation, grasp adaptation through the control of contact motion is demonstrated. Maintaining rolling contact with a moving surface is formulated as an acatastatic system. The proposed scheme involves simultaneously controlling interaction forces as well as the relative (rolling) motion. In all cases, computer simulations results are presented to demonstrate the efficacy of the control schemes.

Characterization of Functionality In A Dynamic Environment
Luca Bogoni
Ruzena Bajcsy
MS-CIS-93-76
GRASP LAB 358

Identifying the functionality in objects means to be able to associate a purpose with them in a specific environment. The purpose depends on the intention of the agent and on the applicability of the object in a particular task. In our investigation of functionality we focus on functionalities which involve changes of physical relation and properties between objects in the environment. A formal model, based on Discrete Event Dynamic System Theory (DEDS), is introduced to define an interactive task for recovering and describing functionality. To observe and control the recovery process we introduce the notion of *piecewise observability* of a task but different sensors. The allows the description of a dynamic system in which neither all events nor the time of their occurrence may be predicted in advance. We have developed an experimental system consisting of actuators and both force and position sensors, for carrying out the interactive recovery of functionality. In particular, we demonstrate how this approach can be used by carrying out some experiments investigating the functionality of piercing. Furthermore, we discuss the importance of a multisensory approach for the observation and interpretation of functionality.

VERSA: A Tool For The Specification and Analysis of Resource-Bound Real-Time Systems
Duncan Clarke
Insup Lee
Hong-liang Xie
MS-CIS-93-77
DISTRIBUTED SYSTEMS LAB 34
GRASP LAB 359

VERSA is a tool that assists in the algebraic analysis of real-time systems. It is based on ACSR, a timed process algebra designed to express resource-bound real-time distributed systems. VERSA is designed to be both a usable and useful tool for the analysis of ACSR specifications. Usability is assured by a flexible user interface that uses ACSR's traditional notation augmented with conventions from programming languages and mathematics that allow concise specification of realistic systems. Usefulness is the result of the breadth of analysis techniques planned and currently implemented, including algebraic terms rewriting and state-spaced exploration based techniques.

A Vision-Based Learning Method for Pushing Manipulation
Marcos Salganicoff
Giorgio Metta
Andrea Oddera
Guilio Sandini
MS-CIS-93-78
GRASP LAB 360

We describe an unsupervised on-line method for learning of manipulative actions that allows a robot to push an object connected to it with a rotational point contact to a desired point in image-space. By observing the results of its actions on the object's orientation in image-space, the system forms a predictive forward empirical model. This acquired model is used on-line for manipulation planning and control as it improves. Rather than explicitly inverting for forward model to achieve trajectory control, a stochastic action selection technique [Moore, 1990] is used to select the most informative and promising actions, thereby integrating active perception and learning by combining on-line improvement, task-directed exploration, and model exploitation. Simulation and experimental results of the approach are presented.

A Direct Approach to Vision Guided Manipulation
Marcos Salganicoff
Giorgio Metta
Andrea Oddera
Giulio Sandini
MS-CIS-93-79
GRASP LAB 361

This paper describes a method for robotic manipulation that uses direct image-space calculation of optical flow information for continuous real-time control of manipulative actions. State variables derived from optical flow measurements are described. The resulting approach is advantageous since it robustifies the system to changes in optical parameters and also simplifies the implementation needed to succeed in the task execution. Two reference tasks and their corresponding experiments are described:

the insertion of a pen into a "cap" (the capping experiment) and the rotational point-contact pushing of an object of unknown shape, mass and friction to a specified goal point in the image-space.

Explicit Forgetting Algorithms for Memory Based Learning
Marcos Salganicoff
MS-CIS-93-80
GRASP LAB 362

Memory-based learning algorithms lack a mechanism for tracking time-varying associative mappings. To widen their applicability, they must incorporate explicit forgetting algorithms to selectively delete observations. We describe Time-Weighted, Locally-Weighted and Performance-Error Weighted forgetting algorithms. These were evaluated with a Nearest-Neighbor Learner in a simple classification task. Locally-Weighted Forgetting outperformed Time-Weighted Forgetting under time-varying sampling distributions and mappings, and did equally well when only the mapping varied. performance-Error forgetting traced about as well as the other algorithms, but was superior since it permitted the Nearest-Neighbor learner to approach the Bayes' misclassification rate when the input/output mapping became stationary.

π-calculus: A Unifying Framework for Programming Paradigms
Srinivas Bangalore
MS-CIS-93-81
LINC LAB 256

π-calculus is a calculus for modeling dynamically changing configurations of a network of communicating agents. this paper studies the suitability of π-calculus as a unifying framework to model the operational semantics of the three paradigms of programming: functional, logic and imperative paradigms. In doing so, the attempt is to demonstrate that π-calculus models a primitive that is pervasive in the three paradigms and to illustrate that the three forms of sequential computing are special instances of concurrent computing.

What's So Special About Kruskal's Theorem and The Ordinal Γ_0? A Survey of Some Results In Proof Theory
Jean H. Gallier
MS-CIS-93-82
LOGIC & COMPUTATION 72

This paper consists primarily of a survey of results of Harvey Friedman about some proof theoretic aspects of various forms of Krusal's tree theorem, and in particular the connection with the ordinal Γ_0. We also include a fairly extensive treatment of normal functions on the countable ordinals, and we give a glimpse of Veblen Hierarchies, some subsystems of second-order logic, slow-growing and fast-growing hierarchies including Girard's result, and Goodstein sequences. The central theme of this paper is a powerful theorem due to Kruskal, the "tree theorem", as well as a "finite miniaturization" of Kruskal's theorem due to Harvey Friedman. These versions of Kruskal's theorem are remarkable from a proof-theoretic point of view because they are *not* provable in relatively strong logical systems. They are examples of so-called "natural independence phenomena", which are considered by more logicians as more natural than the mathematical incompleteness results first discovered by Gödel. Kruskal's tree theorem also plays a fundamental role in computer science, because it is one of the main tools for showing that certain orderings on trees are well founded. These orderings play a crucial role in proving the termination of systems of rewrite rules and the correctness of Knuth-Bandix completion procedures. There is also a close connection between a certain infinite countable ordinal called Γ_0 and Kruskal's theorem. Previous definitions of the function involved in this connection are known to be incorrect, in that, the function is not monotonic. We offer a repaired definition of this function, and explore briefly the consequences of its existence.

A Computational Model of Syntactic Processing: Ambiguity Resolution From Interpretation (Dissertation)
Michael Niv
MS-CIS-93-83
LINC LAB 257

Syntactic ambiguity abounds in natural language, yet humans have not difficulty coping with it. In fact, the process of ambiguity resolution is almost always unconscious. But it is not infallible, however, as example 1 demonstrates.

1. The horse raced past the barn fell.

   This sentence is perfectly grammatical, as is evident when it appears in the following context:

2. Two horses were being shown off to a prospective buyer. One was raced past a meadow and the other was raced past a barn.

Grammatical yet unprocessable sentences such as 1 are called 'garden-path sentences.' Their existence provides

an opportunity to investigate the human sentence processing mechanism by studying how and when it fails. The aim of this thesis is to construct a computational model of language understanding which can predict processing difficulty. The data to be modeled are known examples of garden path and non-garden path sentences, and other results from psycholinguistics.

It is widely believed that there are two distinct loci of computation in sentence processing: syntactic parsing and semantic interpretation. One longstanding controversy is which of these two modules bears responsibility for the immediate resolution of ambiguity. My claim is that it is the latter, and that the syntactic processing module is a very simple device which blindly and faithfully constructs all possible analyses for the sentence up to the current point of processing. The interpretive module services as a filter, occasionally discarding certain of these analyses which it deems less appropriate for the ongoing discourse than their competitors.

This document is divided into three parts. The first is introductory, and reviews a selection of proposals from the sentence processing literature. The second part explores a body of data which has been adduced in support of a theory of structural preferences – one that is inconsistent with the present claim. I show how the current proposal can be specified to account for the available data, and moreover to predict where structural preference theories will go wrong. The third part is a theoretical investigation of how well the proposed architecture can be realized using current conceptions of linguistics competence. In it, I present a parsing algorithm and a meaning-based ambiguity resolution method.

Diagnostic Reasoning and Planning In Exploratory-Corrective Domains (Dissertation)
Ron Rymon
MS-CIS-93-84
LINC LAB 258

I have developed a methodology for knowledge representation and reasoning for agents working in exploratory-corrective domains. Working within the field of Artificial Intelligence in Medicine, I used the specific problem of diagnosis-and-repair in multiple trauma management as both motivation and testbed for my work.

A reasoning architecture is proposed in which specialized diagnostic reasoning and planning components are integrated in a cycle of reasoning and action/perception:

1. A *Goal-Directed Diagnostic* (GDD) reasoner which is predicated on the view that diagnosis is only worthwhile to the extent that it can affect repair decisions and that goals can be used to focus on such. Rather than focusing on a diagnosis object as the primary purpose of the diagnostic process, the GDD reasoner is tasked primarily with generating goals for the planner and with reasoning about whether these goals have been satisfied.

2. A *Progressive Horizon Planner* (PHP) which works by constructing intermediate plans via a combination of plan sketching and selection/optimization subprocesses, and then adapting these plans to reflect new information and goals. For the plan sketching sub-part, I propose a selection-and-ordering planning/scheduling paradigm, taking advantage of the limited interaction between goals.

I have implemented this architecture and reasoning components in TraumAID 2.0 – a consultation system for the trauma management domain. In a blinded comparison, out of 97 real trauma cases, three trauma surgeons have judged management plans proposed by TraumAID 2.0 preferable to the actual care by a ratio of 64:17 and to plans generated by its predecessor TraumAID 1.0 by a ratio of 62:9.

Deterministic Selection on the Mesh and the Hypercube
Sanguthevar Rajasekaran
Shibu Yooseph
MS-CIS-93-85
GRASP LAB 363

In this paper we present efficient deterministic algorithms for selection on the mesh connected computers (referred to as the mesh from hereon) and the hypercube. Our algorithm on the mesh runs in time $O(\frac{n}{p} \log \log p + \sqrt{p} \log n)$ where n is the input size and p is the number of processors. The time bound is significantly better than that of the best existing algorithms when n is large. The run time of our algorithm on the hypercube is $O(\frac{n}{p} \log \log p + T_p^s \log n)$, where T_p^s is the time needed to sort p element on a p-node hypercube. In fact, the same algorithm runs on an network in time $O(\frac{n}{p} \log \log p + T_p^s \log n)$, where T_p^s is the time needed for sorting p keys using p processors (assuming that broadcast and prefix computations take time less than or equal to T_p^s.

A Corpus-Based Approach to Language Learning
(Dissertation)
Eric Brill
MS-CIS-93-86
LINC LAB 259

One goal of computational linguistics is to discover a method for assigning a rich structural annotation to sentences that are presented as simple linear strings of words; meaning can be much more readily extracted from a structurally annotated sentence than from a sentence with no structural information. Also, structure allows for a more in-depth check of the well-formedness of a sentence. There are two phases to assigning these structural annotations: first, a knowledge base is created and second, an algorithm is used to generate a structural annotation for a sentence based upon the facts provided in the knowledge base. Until recently, most knowledge bases were created manually by language experts. These knowledge bases are expensive to create and have not been used effectively in structurally parsing sentences from other than highly restricted domains. The goal of this dissertation is to make significant progress toward designing automate that are able to learn some structural aspects of human language with little human guidance. In particular, we describe a learning algorithm that takes a small structuraally annotated corpus of text and a larger unannotated corpus as input, and automatically learns how to assign accurate structural descriptions to sentences not in the training corpus. The main tool we use to automatically discover structural information about language from corpora is transformation-based error-driven learning. the distribution of errors produced by an imperfect annotator is examined to learn an ordered list of transformations that can be applied to provide an accurate structural annotation. We demonstrate the application of this learning algorithm to part of speech tagging and parsing. Successfully applying this technique to create systems that learn could lead to robust, trainable and accurate natural language processing systems.

Building A Large Annotated Corpus of English: The Penn Treebank
Mitchell P. Marcus
Beatrice Santorini
Mary Ann Marcinkiewicz
MS-CIS-93-87
LINC LAB 260

In this paper, we review our experience with constructing one such large annotated corpus–the Penn Treebank, a corpus consisting of over 4.5 million words of American English. During the first three-year phase of the Penn Treebank Project (1989-1992), this corpus has been annotated for part-of-speech (POS) information. In addition, over half of it has been annotated for skeletal syntactic structure.

A Compiler Project for Translating a C Subset to SPARC Assembly Language
Duncan E. Clarke
Richard P. Paul
MS-CIS-93-89
GRASP LAB 364

We present a complete description of a project for a compiler that translates a subset of the C programming language to SPARC assembler language. The project is suitable for a one semester undergraduate course on compilers and interpreters based on the text of Aho, Sethi, and Ullman, and has been used successfully in that context at the University of Pennsylvania. Output that facilitate scoring, and checkpoints for monitoring the students' progress are integral to the project description.

Robust Hypothesis Testing and Statistical Color Classification
(Dissertation)
Julie a. Adams
MS-CIS-93-90
GRASP LAB 365

The purpose of this research is twofold: (*i*) the development of a mathematical model for statistical color classification; and (*ii*) the testing of this model under controlled conditions.

We consider the following hypothesis testing problem: Let $Z = \theta + V$, where the scalar random variable Z denotes the sampling model, $\theta \in \Omega$ is a location parameter, $\Omega \subset R$, and V is additive noise with cumulative distribution function F. We assume F is uncertain, i.e., $F \in \mathcal{F}$, where $\mathcal{F}$ denotes a given uncertainty class of absolutely continuous distributions with a parametric or semiparametric description. The null hypothesis is $H_0 : \theta \in \Omega,\ F \in \mathcal{F}$ and the alternative hypothesis is $H_1 : \theta \notin \Omega,\ F \in \mathcal{F}$.

Through controlled testing we show that this model may be used to statistically classify colors. The color spectrum we use in these experiments is the Munsell color system which combines the three qualities of color sensation: Hue, Chroma and Value. The experiments show: (*i*) The statistical model can be used to classify colors in the Munsell color system; (*ii*) more robust results are achieved by using a Chroma-Hue match instead of a Perfect match; (*iii*) additional robustness can be achieved by

classifying a color based on measurements averaged over a neighborhood of pixels verses measurements at a single pixel; and (*iv*) a larger color spectrum than the Munsell color system is needed to classify a range of man-made and natural objects.

Proving Properties of Typed λ-Terms Using Realizability, Covers, and Sheves
Jean Gallier
MS-CIS-93-91
LOGIC & COMPUTATION 73

We present a general method for proving properties of typed λ-terms. this method is obtained by introducing a semantic notion of realizability which uses the notion of a cover algebra (as in abstract sheaf theory). For this, we introduce a new class of semantic structures equipped with preorders, called pre-applicative structures. These structures need not be extensional. In this framework, a general realizability theorem can be shown. Kleene's recursive realizability and a variant of Kreisel's modified realizability both fit into this framework. Applying this theorem to the special case of the term model, yields a general theorem for proving properties of typed λ-terms, in particular, strong normalization and confluence. This approach clarifies the reducibility method by showing that the closure conditions on candidates of reducibility can be viewed as sheaf conditions. The above approach is applied to the simply-typed λ-calculus (with types $\rightarrow, x, +$, and $\perp$), and to the second-order (polymorphic λ-calculus (with types $\rightarrow$ and $\forall^2$), for which it yields a new theorem.

Selection and Information: A Class-Based Approach to Lexical Relationships (Dissertation)
Philip S. Resnik
MS-CIS-93-92
LINC LAB 261

Selectional constraints are limitations on the applicability of predicates to arguments. For example, the statement "The number two is blue" may be syntactically well formed, but at some level it is anomalous — BLUE is not a predicate that can be applied to numbers.

In this dissertation, I propose a new, information-theoretic account of selectional constraints. Unlike previous approaches, this proposal requires neither the identification of primitive semantic features nor the formalization of complex inferences based on world knowledge. The proposed model assumes instead that lexical items are organized in a conceptual taxonomy according to class membership, where classes are defined simply as sets — that is, extensionally, rather than in terms of explicit features or properties. Selection is formalized in terms of a probabilistic relationship between predicates and concepts: the selectional behavior of a predicate is modeled as its distributional effect on the conceptual classes of its arguments, expressed using the information-theoretic measure of relative entropy. The use of relative entropy leads to an illuminating interpretation of what selectional constraints are: the *strength* of a predicate's selection for an argument is identified with the quantity of *information* it carries about that argument.

In addition to arguing that the model is empirically adequate, I explore its application to two problems. The first concerns a linguistic question: why some transitive verbs permit implicit direct objects ("John ate $\emptyset$") and others do not ("*John brought $\emptyset$"). It has often been observed informally that the omission of objects is connected to the ease with which the object can be inferred. I have made this observation more formal by positing a relationship between inferability and selectional constraints, and have confirmed the connection between selectional constraints and implicit objects in a set of computational experiments.

Second, I have explored the practical applications of the model in resolving syntactic ambiguity. A number of authors have recently begun investigating the use of corpus-based lexical statistics in automatic parsing; the results of computational experiments using the present model suggest that often lexical relationships are better viewed in terms of underlying conceptual relationships such as selectional preference and concept similarity. Thus the information-theoretic measures proposed here can serve not only as components in a theory of selectional constraints, but also as tools for practical natural language processing.

Understanding Natural Language Instructions: A Computational Approach to Purpose Clauses (Dissertation)
Barbara Di Eugenio
MS-CIS-93-93
LINC LAB 262

Human agents are extremely flexible in dealing with Natural Language instructions. I argue that most instructions don't exactly mirror the agent's't knowledge, but are understood by *accommodating* them in the context of the general plan the agent is considering; the accommodation process is guided by the *goal(s)* that the agent is trying to achieve. Therefore a NL system which interprets instructions must be able to recognize and/or hypothesize goals; it must make use of a flexible knowledge represen-

tation system, able to support the specialized inferences necessary to deal with input action description that do not exactly match the stored knowledge.

The data that support my claim are Purpose Clauses (PC's), infinitival constructions as in *Do α to do β*, and Negative Imperatives. I present a pragmatic analysis of both PCs and negative Imperatives. Furthermore, I analyze the computational consequences of PCs, in terms of the relations between actions PCs express, and of the inferences an agent has to perform to understand PCs.

I propose an action representation formalism that provides the required flexibility. It has two components. The *Terminological Box* (TBox) encodes *linguistic* knowledge about actions, and is expressed by means of the hybrid system CLASSIC [Brachman et al., 1991].

To guarantee that the primitives of the representation are linguistically motivated, I derive them from Jackendoff's work on Conceptual Structures [1983; 1990]. The Action Library encodes *planning* knowledge about actions. The action terms used in the plans are those defined in the TBox.

Finally, I present an algorithm that implements inferences necessary to understand *Do α to do β*, and supported by the formalism I propose. In particular, I show how the TBox classifier is used to infer whether α can be assumed to match one of the substeps in the plan for β, and how expectations necessary for the match to hold are computed.

Facilitating Transformations in a Human Genome Project Database
S.B. Davidson
A.S. Kosky
B. Eckman
MS-CIS-93-94
LOGIC & COMPUTATION 74

Human Genome Project databases present a confluence of interesting database challenges: rapid schema and data evolution, complex data entry and constraint management, and the need to integrate multiple data sources and software systems which range over a wide variety of models and formats. While these challenges are not necessarily unique to biological databases, their combination, intensity and complexity are unusual and make automated solutions imperative. We illustrate these problems in the context of the Human Genome Database for Chromosome 22 (Chr22DB), and describe a new approach to a solution for these problems, by means of a deductive language for expressing database transformations and constraints.

A Bounded Degree Property and Finite-Cofiniteness of Graph Queries
Leonid Libkin
Limsoon Wong
MS-CIS-93-95
LOGIC & COMPUTATION 75

We provide new techniques for the analysis of the expressive power of query languages for nested collections. These languages may use set or bag semantics and may be further complicated by the presence of aggregate functions. We exhibit certain classes of graphics and prove that properties of these graphics that can be tested in such languages are either finite or cofinite. This result settles that conjectures of Grumbach, Milo, and Paredaens that parity test, transitive closure, and balanced binary tree test are not expressible in bah languages like BALG of Grumbach and Milo and $\mathcal{BQL}$ of Libkin and Wong. Moreover, it implies that many recursive queries, including simple ones like test for a chain, cannot be expressed in a nested relational language even when aggregate functions are available. In an attempt to generalize the finite-cofiniteness result, we study the bounded degree property which says that the number of distinct in- and out-degrees in the output of a graph query does not depend on the size of the input if the input is "simple." We show that such a property implies a number of inexpressibility results in a uniform fashion. We then prove the bounded degree property for the nested relational language.

Extended Intensity Range Imaging
Brian C. Madden
MS-CIS-93-96
GRASP LAB 366

A single composite image with an extended intensive range is generated by combining disjoining regions from different images of the same scene. The set of images is obtained with a charge-couple device (CCD) set for different flux integration times. By limiting differences in the integration times so that the ranges of output pixel values overlap considerably, individual pixels are assigned the value measured at each spatial location that is in the most sensitive range where the values are both below saturation and are most precisely specified. Integration times are lengthened geometrically from a minimum where all pixel values are below saturation until all dark regions emerge from the lowest quantization level. the method is applied to an example scene and the effect the composite images have on traditional low-level imaging methods also is examined.

Convex Hulls: Complexity and Applications (A Survey)
Suneeta Ramaswami
MS-CIS-93-97

Computational geometry is, in brief, the study of algorithms for geometric problems. Classical study of geometry and geometric objects, however, is not well-suited to efficient algorithms techniques. thus, for the given geometric problems, it becomes necessary to identify properties and concepts that lend themselves to efficient computation. The primary focus of this paper will be on one such geometric problems, the Convex Hull problem.

Algorithmic Motion Planning and Related Geometric Problems on Parallel Machines (Dissertation Proposal)
Suneeta Ramaswami
MS-CIS-93-98

Optimal Parallel Randomized Algorithms for the Boronoi Diagram of Line Segments In The Plane and Related Problems
Sanguthevar Rajasekaran
Suneeta Ramaswami
MS-CIS-93-99

In this paper, we present an optimal parallel randomized algorithm for the Voronoi diagram of a set of n non-intersecting (except possibly at endpoints) line segments in the plane. Our algorithm runs in $O(\log n)$ time with very high probability and uses $O(n)$ processors on a CRCW PRAM. This algorithm is optimal in terms of $P.T$ bounds since the sequential time bound for this problem is $\Omega(n \log n)$. Our algorithm improves by an $O(\log n)$ factor the previously best known deterministic parallel algorithm which runs in $O(\log^2 n)$ time using $O(n)$ processors. We obtains this result by using random sampling at "two stages" of our algorithm and using efficient randomized search techniques. This technique gives a direct optimal algorithm for the Voronoi diagram of points as well (all other optimal parallel algorithms for this problem use reduction from the 3-d convex hull construction).

Control of Visually Guided Behaviors
Jana Kosecká
Ruzena Bajcsy
Max Mintz
MS-CIS-93-101
GRASP LAB 367

We propose an approach for modeling visually guided behaviors of agents which explore and navigate in unknown and partially known environments. Behaviors are modeled as finite state machines (FSM), where the states of the model correspond to particular continuous control strategies and the transitions between them are caused by events representing qualitative or asynchronous changes in the behavior evolution. In order to prevent conflicts in parallel execution of multiple behaviors we adopt the supervisory control theory of discrete Event System (DES). Modeling the participating processes using the DES framework allows us to capture often complex interactions between components of the system and synthesize the resulting supervisor, guaranteeing the overall controllability of the system at the discrete event level. In the real world agents have multiple options/paths for carrying out their task. Hence there is a need for selecting different control strategies based on efficiency and safety criteria. We have included in our formalism a measure of efficiency as the nominal cost (in our case, the traversal time) and a measure of safety at the risk cost (in our case, the inverse of the distance between the agent and obstacles). Experiments have been carried out testing the described formalism with one agent carrying out the task of avoiding an obstacle in its path while tracking a target.

Using Context To Specify Intonation In Speech Synthesis
Scott Prevost
Mark Steedman
MS-CIS-93-102
LINC LAB 263

A generator based on Combinatory Categorial Grammar using a simple and domain-independent discourse model can be used to direct synthesis of intonation contours for responses to data-base queries, conveying distinctions of contrast and emphasis determined by the discourse model and the state of the knowledge-base.

Spherical Retinal Flow for A Fixating Observer

Inigo Thomas
Eero Simoncelli
Ruzena Bajcsy

MS-CIS-94-06
GRASP LAB 368

When a human observer moves, the eye continually fixates on targets in the world. Although fixation is a common process in human vision, its role has not yet been established for computational purposes. The main contribution of this paper is the formalize the retinal flow for a fixating observer. A further contribution – a potentially more practical one – is to explore the role of the periphery in predicting collision. Utilizing fixation is expected to turn out be especially fruitful in light of recent advance in computer vision for constructing active head/eye systems.

In this work we make the following assumption: (i) the observer moves with respect to the world and fixates on a target; (ii) the world is rigid, with no independently moving elements; and (iii) the possible rotation axes of the eye lie on a plane (comparable to Listing's Plane). Assumption (ii) and (iii) make the problem of determine retinal flow tractable. We first define retinal flow for a 2D universe and then extend it to the full 3D will be decomposed into longitudinal and latitudinal flow; the behavior of longitudinal flow along the retinal periphery will be further analyzed for interesting properties. Finally the results of a simulated experiment on retinal flow at the periphery will be presented.

Massively Parallel Simulation of Structured Connectionist Networks: An Interim Report

D.R. Mani
Lokendra Shastri

MS-CIS-94-10
LINC LAB 264

We map structured connectionist models of knowledge representation and reasoning onto existing general purpose massively parallel architectures with the objective of developing and implementing practical, real-time knowledge base systems. SHRUTI, a connectionist knowledge representation and reasoning system which attempts to model reflexive reasoning, will serve as our representative connectionist model. Efficient simulation systems for SHRUTI are developed on the Connection Machine CM-2—an SIMD architecture—and on the Connection Machine CM-5—a MIMD architecture. The resulting simulators are evaluated and tested using large, random knowledge bases with up to half a million rules and facts.

Though SIMD simulations on the CM-2 are reasonably fast—requiring a few seconds to tens of seconds for answering simple queries—experiments indicate that MIMD simulations are vastly superior to SIMD simulations and offer hundred- to thousand-fold speedups.

This work provides new insights into the simulation of structured connectionist networks on massively parallel machines and is a step toward developing large yet efficient knowledge representation and reasoning systems.

Efficient Compilation of High-Level Data Parallel Algorithms

Dan Suciu
Val Tannen

MS-CIS-94-17
LOGIC & COMPUTATION LAB 78

We present a high-level parallel calculus for nested sequences, NSC, offered as a possible theoretical "core" of an entire class of collection-oriented parallel languages. NSC is based on while-loops as opposed to general recursion. A formal, machine independent definition of the parallel time complexity and the work complexity of programs in NSC is given. Our main results are: (1) We give a translation method for a particular form of recursion, called map-recursion, into NSC, that preserves the time complexity and adds an arbitrarily small overhead to the work complexity, and (2) We give a compilation method for NSC into a very simple vector parallel machine, which preserves the time complexity and again adds an arbitrarily small overhead to the work complexity.

View Selection Strategies for Multi-View, Wide-Baseline Stereo

Hany Farid
Sang Wook Lee
Ruzena Bajcsy

MS-CIS-94-18
GRASP LAB 373

Recovering 3D depth information from two or more 2D intensity images is a long standing problem in the computer vision community. This paper presents a multi-baseline, coarse-to-fine stereo algorithm which utilizes any number of images (more than 2) and multiple image scales to recover 3D depth information. Several "view-selection strategies" are introduced for combining information across the multi-baseline and across scale spade. The control strategies allow us to exploit, maximally, the benefits of large and small baselines and mask sizes while minimizing errors. Results of recovering 3D depth information from a human head are presented. The resulting

depth maps are of good accuracy with a depth resolution of approximately 5mm.

The Well-tempered Computer
Mark Steedman
MS-CIS-94-19 LINC LAB 266

The psychological mechanism by which even musically untutored people can comprehend novel melodies resembles that by which they comprehend sentences of their native language. the paper identifies a syntax, a semantics, and a domain or "model". These elements are examined in application to the task of harmonic comprehension and analysis of unaccompanied melody, and a computational theory is argued for.

Process Algebra, CCS, and Bisimulation Decidability
Seth Kulick
MS-CIS-94-20
LINC LAB 267

Over the past fifteen years, there has been intensive study of formal systems that can model concurrency and communication. Two such systems are the Calculus of Communicating Systems, and the Algebra of Communicatting Process. The objective of this paper has two aspects: (1) to study the characteristics and features of these two systems, and (2) to investigate two interesting formal proofs concerning issues of decidability of bisimulation equivalence in these systems. An examination of the processes that generate context-free languages as a trace set shows that their bisimulation equivalence is decidable, in contrast to the undecidability of their trace set equivalence problem for processes with a limited amount of concurrency is decidable.

Approximation in Databases
Leonid Libkin
MS-CIS-94-21
LOGIC & COMPUTATION 79

One source of partial information in databases is the need to combine information from several databases. Even if each database is complete for some "world", the combined databases will not be, and answers to queries against such combined databases can only be approximated. In this paper we describe various situations in which a precise answer cannot be obtained for a query asked against multiple databases. Based on an analysis of these situations, we propose a classification of constructs that can be used to model approximations.

One of the main goals is to show that most of these models of approximations possess universality properties. The main motivation for doing this is applying the data-oriented approach, which turns universality properties into syntax, to obtain languages for approximations. We show that the languages arising from the universality properties have a number of limitations. In an attempt to overcome those limitations, we explain how all the languages can be embedded into a language for conjunctive and disjunctive sets from [21], and demonstrate its usefulness in querying independent databases.

State Minimization for Concurrent System Analysis Based on State Space Exploration
Inhye Kang
Insup Lee
MS-CIS-94-22
LOGIC & COMPUTATION 80

A fundamental issue in the automated analysis of concurrent systems is the efficient generation of the reachable state space. Since it is not possible to explore all the reachable states of a system if the number of states is very large or infinite, we need to develop techniques for minimizing the state space. this paper presents our approach to cluster subsets of states into equivalent classes. We assume that concurrent systems are specified as communicating state machines with arbitrary data space. We describe a procedure for constructing a minimal reachability state graph from communicating state machines. As an illustration of our approach, we analyze a producer-consumer program written in Ada.

Modeling the Interaction between Speech and Gesture
Justine Cassell
Matthew Stone
Brett Douville
Scott Prevost
Brett Achorn
Mark Steedman
Norm Badler
Catherine Pelachaud
MS-CIS-94-23
LINC LAB 268 GRAPHICS LAB 61

This paper describes an implemented system that generates spoken dialogue, including speech, intonation, and gesture, using two copies of an identical program that differ only in knowledge of the world and which must cooperate to accomplish a goal. The output of the dialogue generation is used to drive a three-dimensional interactive animated model – two graphic figures on a computer screen who speak and gesture according to the rules of the system. The system is based upon a formal, predictive

and explanatory theory of the gesture-speech relationship. A felicitous outcome is a working system to realize autonomous animated conversational agents for virtual reality and other purposes, and a tool for investigating the relationship between speech and gesture.

A Process Algebra of Communicating Shared Resources with Dense Time and Priorities (Dissertation)
Patrice Brémond-Grégoire
MS-CIS-94-24
LOGIC & COMPUTATION 81

The correctness of real-time distributed systems depends not only on the function they compute but also on their timing characteristics. Furthermore, those characteristics are strongly influenced by the delays due to synchronization and resource availability. Process algebras have been used successfully to define and prove correctness of distributed systems. More recently, there has been a lot of activity to extend their application to real-time systems. The problem with most current approaches is that they ignore resource constraints and assume either a total parallelism (unlimited resources) or total interleaving (single resource.) Algebra of Communicating Shared Resources (ACSR) is a process algebra designed for the formal specification and manipulation of distributed systems with resource and real-time constraints. A dense time domain provides a more natural way of specifying systems compared to the usual discrete time. Priorities provide a measure of urgency for each action and can be used to ensure that deadlines are met. In ACSR, processes are specified using resource bound, timed actions and instantaneous synchronization events. Processes can be combined using traditional operators such as nondeterministic choice and parallel execution. Specialized operators allow the specification of real-time behavior and constraints. The semantics of ACSR is defined as a labeled transition system. Equivalence between processes is based on the notion of strong bisimulation. A sound and complete set of algebraic laws can be used to transform also any ACSR process into a normal form. In practice, several specifications may satisfy the same requirements with various degree of desirability. Some may use more resources; some may be faster. In fact, there are many ways to rank processes. We describe a method for defining order relations between execution traces and further expanding the relation to general processes. Monotonicity is an important property of operators as it ensures that ordering is preserved by contexts. We study the conditions that must be satisfied by the trace ordering to ensure monotonicity at the process level, both in the prioritized and unprioritized cases. While most operations are monotonic for a large variety of trace relations, few retain this property in a prioritized setting.

Automatic Synthesis of Simplified 3D Models From Detailed Data (Dissertation)
Eunyoung Koh
MS-CIS-94-25
GRAPHICS LAB 62

Modeling and displaying complex objects in Computer Graphics often requires using multiple levels of detail for display practicality and efficiency. We

Natural Language Processing
Mark Steedman
MS-CIS-94-32
LINC LAB 270

The subject of Natural Language Processing can be considered in both broad and narrow senses. In the broad sense, it covers processing issues at all levels of natural language understanding, including speech recognition, syntactic and semantic analysis of sentences, reference to the discourse context (including anaphora, inference of referents, and more extended relations of discourse coherence and narrative structure), conversational inference and implicature, and discourse planning and generation. In the narrower sense, it covers the syntactic and semantic processing *sentences* to deliver semantic objects suitable for referring, inferring, and the like. Of course, the results of inference and reference may under some circumstances play a part in processing in the narrow sense. But the processes that are characteristic of these other modules are not the primary concern.

Critiquing: Effective Decision Support In Time-Critical Domains (Dissertation Proposal)
Abigail S. Gertner
MS-CIS-94-35
LINC LAB 272

The effective communication of information is an important concern in the design of an expert consultation system. Several researchers have chosen to adopt a *critiquing* mode, in which the system evaluates8 and reacts to a solution proposed by the user rather than presenting its own solution. In this proposal, I present an architecture for a critiquing system that functions in real-time, *during* the process of developing and executing a management plan in time-critical situations. the architecture is able to take account of and reason about multiple, interacting goals and to identify critical errors in the proposed management plan. This architecture is being implemented as part of the TraumAID system for the management of patients with severe injuries.

An Efficient Generation of the Timed Reachability Graph for the Analysis of Real-Time Systems
Inhye Kang
Insup Lee
MS-CIS-94-36
DISTRIBUTED SYSTEMS LAB 79
LOGIC & COMPUTATION LAB 83

As computers become ubiquitous, they are increasingly used in safety critical environments. Since many safety critical applications are real-time systems, automated analysis technique of real-time properties is desirable. Most widely used automated analysis techniques are based on state space exploration. Automatic analysis techniques based on state space exploration suffer from the state space explosion problem. In particular, a real-time system may have an unbounded number of states due4 to infinitely many possible time values. This paper presents our approach for generating a finite and efficient representation of the reachable states called a timed reachability graph for a real-time system is specified using a timed automaton which is a timed extension of the well-known finite automaton. Our approach for coping with the state explosion problem is to extract timing information from states and to represent it as relative time relations between transitions. We also present an algorithm for computing the minimum and maximum time bounds between executions of two actions from a timed reachability graph to determine timing properties.

Specifying Intonation from Context for Speech Synthesis
Scott Prevost
Mark Steedman
MS-CIS-94-37
LINC LAB 273

This paper presents a theory and a computational implementation for generating prosodically appropriate synthetic speech in response to database queries. Proper distinctions of contrast and emphasis are expressed in an intonation contour that is synthesized by rule under the control of a grammar, a discourse model, and a knowledge base. The theory is based on Combinatory Categorial Grammar, a formalism which easily integrates the notions of syntactic constituency, semantics, prosodic phrasing and information structure. Results from our current implementation demonstrate the system's ability to generate a variety of intonational possibilities for a given sentence depending on the discourse context.

Centering: A Framework for Modelling the Coherence of Discourse
Barbara J. Grosz
Aravid K. Joshi
Scott Weinstein
MS-CIS-94-40
LINC LAB 274

Our original paper (Grosz, Joshi, and Weinstein, 1983) on centering claimed that certain entities mentioned in an utterance were more central than others and that this property imposed constraints on a speaker's use of different types of referring expression. Centering was proposed as a model that accounted for this phenomenon. We argued that the compatibility of centering properties of an utterance with choice of referring expression affected the coherence of discourse. Subsequently, we expanded the ideas presented therein. We defined various centering constructs and proposed two centering rules in terms of these constructs. A draft manuscript describing this elaborated centering framework and presenting some initial theoretical claims has been in wide circulation since 1986. this draft (Grosz, Joshi, and Weinstein 1986, hereafter, **GJW86**) has led to a number of papers by others on this topic and has been extensively cited, but has never been published.

We have been urged to publish the more detailed description of the centering framework and theory proposed in **GJW86** so that an official version would be archivally available. The task of completing and revising this draft became more daunting as time passed and more and more papers appeared on centering. Many of these papers pro-

posed extensions to or revisions of the theory and attempted to answer questions posed in **GJW86**. It has become ever more clear that it would be useful to have a "definitive" statement of the original motivations for centering, the basic definitions underlying the centering framework, and the original theoretical claims. This paper attempts to meet that need. To accomplish this goal, we have chosen to removed descriptions of many open research questions posed in **GJW86** as well as solutions that were only partially developed. We have also greatly shortened the discussion of criteria for and constraints on a possible semantic theory as a foundation for this work.

Jack/TTES: A System for Production and Real-time Playback of Human Figure Motion in a DIS Environment

John P. Granieri
MS-CIS-94-42
HUMAN MODELING & SIMULATION LAB 65

This document describes a modified *Jack* system for off-line motion production and on-line (real-time) motion playback to an external IRIS-Performer-based host rendering system. This work was done in partial fulfillment of Contract #N61339-94-C-0005 for the US Marine Corps through NAWCTSD (Naval Air Warfare Center, Training systems Division).

The work described herein was contributed by several of the members of the Center for Human Modeling and Simulation: John Granieri (Design/Engineering/Integration), Rama Bindiganavale (animator, posture transitions), Hanns-Oskar Poor (animator, posture transitions, Hyeongseok Ko (walking and running motion), Micheal Hollick (locomotion playback control), Bond-Jay Ting (body sculpting), Francisco Azoula (body sculptin, anthropometry), Pei-Hwa Ho (body normalization), Jonathan Crabtree (Performer, TIPS file format), Xinmin Zhao (slaving), Zhongyang Feng (DIS logfile player), Welton Becket and Barry Reich (terrain reasoning and reactive agent control).

Behavior-Based Control for Time-Delayed Teleoperation (Dissertation)

Matthew R. Stein
MS-CIS-94-43
GRASP LAB 378

Remote control of robotic manipulation has applications in undersea, shallow space and low bandwidth communication environments. Communication delays on the order of several seconds can occur during space operations that use relay stations, undersea operations using acoustic communication channels, or when significant computational delay exists. This system is also applicable to time varying time delays as experienced when the Internet is used as the medium of communication. The use of supervisory control to address the time delay problem requires a remote manipulator to exhibit some degree of autonomy. This autonomy is necessary to minimize the contribution of communication delay time to task completion time and to relieve the operator of the responsibility for contact control. This dissertation examines the use of behavior-based or subsumption architecture control to provide the required autonomy for the remote manipulator. Behavior-based controllers demonstrate desirable features including reliability and robust operation in unstructured environments and, when used in conjunction with operator direction, avoids the need for higher level representations which do not fit well into the architecture. We develop a model for communications between the operator and the behavior-based controller and define the human-machine interface for operator direction. We demonstrate the supervisory control system on a GRASP Laboratory mockup of a slicing task performed during a satellite repair operation. In this task a robot cuts securing tape along the seams of the panels of a thermal protection blanket. We tested the validity of the supervisory control system by performing controlled experiments using untrained operators. Quantitative and qualitative results demonstrate the supervisory control concept is a practical and viable solution to the time delay problem.

Implementing Selective Attention in Machines: The Case of Touch-Driven Saccades

Michele Rucci
Ruzena Bajcsy
MS-CIS-94-44
GRASP LAB 379

Recent paradigms in the fields of robotics and machine perception have emphasized the importance of selective attention mechanisms for perceiving and interacting with the environment. In the case of a system involved in operations requiring a physical interaction with the surrounding environment, a major role is played by the capability of attentively responding to tactile events. By performing somatosensory saccades, the nature of the cutaneous stimulation can be assessed, and new motor actions can be planned. However, the study of touch–driven attention, has almost been neglected by robotics researchers. In this paper the development of visuo-cutaneo coordination for the production of somatosensory saccades is investigated, and a general architecture for integrating different kinds of attentive mechanisms is proposed. The system autonomously discovers the sensorymotor trans-

formation which links tactile events to visual saccades, on the basis of multisensory consistencies and basic, built–in, motor reflexes. Results obtained both with simulations and robotic experiments are analyzed.

Trajectory Formulation In Human Movement: An Extension of Existing Models For Single-Arm Motions To Coupled Motions Of Two Arms (Dissertation)
Gregory J. Garvin
MS-CIS-94-45
GRASP LAB 380

Although there have been many studies of the kinematics and dynamics of single-arm motions little has been done in this area with regards to motions involving multiple limbs, especially for the case of dynamic interaction between two arms. The primary goal of this research has been to examine the manipulation of an object requiring the constant use of both arms. Specifically, two extensions of the minimum jerk model [Fla85], which has been successful in past studies of the single-arm case, were developed to account for the case of two arms manipulating a planar object while dynamically coupled. In order to quantitatively study two-arm manipulation by humans, a planar, three degree of freedom manipulandum capable of measuring the positions of and forces at the hands was designed and built. The end-effector of the manipulandum consisted of two handles mounted at opposite ends of an aluminum bar. Volunteer subjects participated in experiments in which they were asked to move between a series of illuminated targets which indicated the desired position and orientation of the bar. Although the proposed models did successfully predict several aspects of the trajectories, they were unable to account for others. However, it may be possible to modify the models to account for the observed behavior by relaxing some of the simplifying assumptions made to facilitate their development. Furthermore, the analysis of the empirical data has also suggested several alternative approaches to modeling two-arm motion.

Description Based Parsing In A Connectionist Network (Dissertation)
James B. Henderson
MS-CIS-94-46
LINC LAB 275

Recent developments in connectionist architectures for symbolic computation have made it possible to investigate parsing in a connectionist network while still taking advantage of the large body of work on parsing in symbolic frameworks. This dissertation investigates syntactic parsing in the temporal synchrony variable binding model of symbolic computation in a connectionist network. This computational architecture solves the basic problem with previous connectionist architectures, while keeping their advantages. However, the architecture does have some limitations, which impose computational constraints on parsing in this architecture. This dissertation argues that, despite these constraints, the architecture is computationally adequate for syntactic parsing, and that these constraints make significant linguistic predictions. To make these arguments, the nature of the architecture's limitations are first characterized as a set of constraints on symbolic computation. This allows the investigation of the feasibility and implications of parsing in the architecture to be investigated at the same level of abstraction as virtually all other investigations of syntactic parsing. Then a specific parsing model is developed and implemented in the architecture. The extensive use of partial descriptions of phrase structure trees is crucial to the ability of this model to recover the syntactic structure of sentences within the constraints. Finally, this parsing model is tested on those phenomena which are of particular concern given the constraints, and on an approximately unbiased sample of sentences to check for unforeseen difficulties. The results show that this connectionist architecture is powerful enough for syntactic parsing. They also show that some linguistic phenomena are predicted by the limitations of this architecture. In particular, explanations are given for many cases of unacceptable center embedding, and for several significant constraints on long distance dependencies. These results give evidence for the cognitive significance of this computational architecture and parsing model. This work also shows how the advantages of both connectionist and symbolic techniques can be unified in natural language processing applications. By analyzing how low level biological and computational considerations influence higher level processing, this work has furthered our understanding of the nature of language and how it can be efficiently and effectively processed.

Building and Control In CCG and Its Relatives
Mark Steedman
MS-CIS-94-47
LINC LAB 276

The CCG account of the unbounded constructions – in particular, relativisation and coordination – generalizes the notion of surface structure in a way that disrupts traditional notions of dominance and command. This has led

researchers in other frameworks to suggest that the theory is fundamentally incompatible with a coherent theory of binding and control – the bounded constructions. The present paper offers a theory of binding in CCG which preserves the original account of the unbounded dependencies, and which renders it immediately compatible with other theories, TAG in particular. The theory requires the abandonment of one assumption that has been traditional (though not essential) in other categorial approaches. The significance of this move is discussed.

Domain-Independent Queries on Databases with External Functions
Dan Suciu
MS-CIS-94-48
LOGIC & COMPUTATION 86

We investigate queries in the presence of external functions with arbitrary inputs and outputs (atomic values, sets, nested sets etc). We propose a new notion of domain independence for queries with external functions which, in contrast to previous work, can also be applied to query languages with fixpoints or other kinds of iterators. Next, we define two new notions of computable queries with external functions, and prove that they are equivalent, under the assumption that the external functions are total. Thus, our definition of computable queries with external functions is robust. Finally, based on the equivalence result, we give examples of complete query languages with external functions. A byproduct of the equivalence result is the fact that Relational Machines are complete for complex objects: it was known that they are not complete over flat relations.

Deformable Models with Parameter Functions: Application to Heart-Wall Modeling
Jinah Park
Dimitri Metaxas
Alistair Young
MS-CIS-94-49
HUMAN MODELING & SIMULATION LAB 66

This paper develops a new class of physics-based deformable models which can deform both globally and locally. their global parameters are functions allowing the definition of new parameterized primitives and parameterized global deformations. These new global parameter functions improve the accuracy of shape description through the use of a few intuitive parameters such as functional bending and twisting. Using a physics-based approach we convert these geometric models into deformable models that deform due to forces exerted from the datapoints so as to conform to the given dataset. We present an experiment involving the extraction of shape and motion of the Left Ventricle (LV) of a heart from MRI-SPAMM data based on a few global parameter functions.

Model-based Analysis of Cardiac Motion from Tagged MRI Data
Jinah Park
Dimitri Metaxas
Alistair Young
Leon Axel
MS-CIS-94-50
HUMAN MODELING & SIMULATION LAB 67

We develop a new method for analyzing the motion of the left ventricle (LV) of a heart from tagged MRI data. Our technique is based on the development of a new class of physics-based deformable models whose parameters are functions allowing the definition of new parameterized primitives and parameterized deformations. These parameter functions improve the accuracy of shape description through the use of a few intuitive parameters such as functional twisting. Furthermore, these parameters require no complex post-processing in order to be used by a physician. Using a physics-based approach, we convert these geometric models into deformable models that deform due to forces exerted from the datapoints and conform to the given dataset. We present experiments involving the extraction of shape and motion of the LV from MRI-SPAMM data based on a few parameter functions. Furthermore, by plotting the variations over time of the extracted model parameters from normal and abnormal heart data we are able to characterize quantitatively their-differences.

A Multiple-Conclusion Meta-Logic
Dale Miller
MS-CIS-94-51
LINC LAB 277

The theory of cut-free sequent proofs has been used to motivate and justify the design of a number of logic programming languages. Two such languages, λProlog and its linear logic refinement, Lolli, provide for various forms of abstraction (modules, abstract data types, higher-order programming) but lack primitives for concurrency. The logic programming language, LO (Linear Objects) provides for concurrency but lacks abstraction mechanisms. In this paper we present Forum, a logic programming presentation of all of linear logic that modularly extends the languages λProlog, Lolli, and LO. Forum, therefore, allows specifications to incorporate both abstractions and concurrency. As a meta-language, Forum greatly extends

the expressiveness of these other logic programming languages. To illustrate its expressive strength, we specify in Forum a sequent calculus proof system and the operational semantics of a functional programming language that incorporates such nonfunctional features as counters and references.

Formal and Computational Aspects of Natural Language Syntax (Dissertation)
Owen Rambow
MS-CIS-94-52
LINC LAB 278

This thesis explores issues related to using a restricted mathematical formalism as the formal basis for the representation of syntactic competence and the modeling of performance. The specific contribution of this thesis is to examine a language with considerable freer word-order than English, name German, and to investigate the formal requirements that this syntactic freedom imposes. Free word order (or free constituent order) languages can be seen as a test case for linguistic theories, since presumably the stricter word order can be subsumed by an apparatus that accounts for freer word order.

The formal systems investigated in this thesis are based on the tree adjoining grammar (TAG) formalism of Joshi et al. (1975). TAG is an appealing formalism for the representation of natural language syntax because its elementary structures are phrase structure trees, which allows the linguist to localized linguistic dependencies such as agreement, subcategorization, and filler-gap relations, and to develop a theory of grammar based on the lexicon.

The main results of the thesis are an argument that simple TAGs are formally inadequate, and the definition of an extension of TAG that is. Every aspect of the definition of this extension to TAG, called V-TAG, is specifically motivated by linguistic facts, not by formal considerations. A formal investigation of V-TAG reveals that (when lexicalized) it has restricted generative capacity, that it is polynomial parsable, and that it forms an abstract family of languages. This means that is has desirable formal properties for representing natural language syntax. both a formal automaton and a parser for V-TAG are presented.

As a consequence of the new system, a reformulation of the linguistic theory that has been proposed for TAG suggests itself. Instead of including a transformational step in the theory of grammar, all derivations are performed within mathematically defined formalisms, thus limiting the degrees of freedom in the linguistic theory, and making the theory more appealing from a computational point of view. This has several interesting linguistic consequences; for instance, functional categories are expressed by feature content (not node labels), and head movement is replaced by the adjunction of heads. The thesis sketches a fragment of a grammar of German, which covers phenomena such as scrambling, extraposition, topicalization, and the V2 effect.

Finally, the formal automaton for V-TAG is used as a model of human syntactic processing. It is shown that this model makes several interesting predictions related to free word order in German.

A Computational Approach to Aspectual Composition (Dissertation)
Michael White
MS-CIS-94-53
LINC LAB 279

In recent years, it has become common in the linguistics and philosophy literature to assume that events and processes are ontologically distinct entities, on a par with objects and substances. At the same time, the idea that time-based (episodic) knowledge should be represented as a collection of interrelated eventualities has gained increasing acceptance in the computational linguistics and artificial intelligence literature.

Contrary to what one might expect, a search through the prior literature in linguistics and philosophy reveals no account in which these sortal distinctions play a direct role in adequately explaining the problem of aspectual composition and the closely related imperfective paradox. In the computational linguistics and artificial intelligence literature, moreover, relatively little attention has been paid to either problem.

In the first part of the dissertation, I investigate the hypothesis that the parallel ontological distinctions introduced above may be directly employed in an explanatory formal account of the problem of aspectual composition and the imperfective paradox. In so doing, I develop a synthesis of proposals by Hinrichs (1985), Krifka (1989; 1992) and Jackendoff (1991) which makes correct predictions in many cases not considered by these authors. In particular, the account is the first to adequately explain the syntactic and semantic behavior of non-individuating accomplishment expressions, such as *Jack pour some amount of wort into the carboy*, which are to vague to individuate a single event by nevertheless behave like other Vendlerian accomplishments.

In the second part of the dissertation, I explore the potential computational applications of the linguistic account, by way of two case studies. In the first one, I follow Moens (1987) in showing how a calculus of eventualities can facilitate the implementation of a simple statement

verifier which allows for a much greater range of natural language queries than is usually the case with temporal databases. In the second, more preliminary study, I examine the relevance of the model-theoretic analysis to discourse interpretation, within the context of devising a program which produces simple microworld animations using short narrative descriptions as input specifications.

Querying Nested Collections (Dissertation)
Limsoon Wong
MS-CIS-94-54
LOGIC & COMPUTATION 87

This dissertation investigates a new approach to query languages inspired by structural recursion and by the categorical notion of monad.

A language based on these principles has been designed and studied. It is found to have the strength of several widely known relational languages but without their weaknesses. This language and its various extensions are shown to exhibit a conservative extension property, which indicates that the depth of nesting of collections in intermediate data has no effect on their expressive power. These languages also exhibit the finite-cofiniteness property on many classes of queries. These two properties provide easy answers to several hitherto unresolved conjectures on query languages that are more realistic than the flat relational algebra.

A useful rewrite system has been derived from the equational theory of monads. It forms the core of a source-to-source optimizer capable of performing filter promotion, code motion, and loop fusion. Scanning routines and printing routines are considered as part of optimization process. An operational semantics that is a blending of eager evaluation and lazy evaluation is suggested in conjunction with these input-output routines. This strategy leads to a reduction in space consumption and a faster response time while preserving good total time performance. Additional optimization rules have been systematically introduced to cache and index small relations, to map monad operations to several classical join operators, to cache large intermediate relations, and to push monad operations to external servers.

A query system Kleisli and a high-level query language CPL for it have been built on top of the functional language ML. Many of my theoretical and practical contributions have been physically realized in Kleisli and CPL. In addition, I have explored the idea of open system in my implementation. Dynamic extension of the system with new primitives, cost functions, optimization rules, scanners, and writers are fully supported. As a consequence, my system can be easily connected to external data sources. In particular, it has been successfully applied to integrate several genetic data sources which include relational databases, structured files, as well as data generated by special application programs.

Final Report to NSF of the Standards for Facial Animation Workshop
Catherine Pelachaud
Norman I. Badler
Marie-Luce Viaud
MS-CIS-94-56
HUMAN MODELING & SIMULATION LAB 68

The human face is an important and complex communication channel. It is a very familiar and sensitive object of human perception. The facial animation field has increased greatly in the past few years as fast computer graphics workstations have made the modeling and real-time animation of hundreds of thousands of polygons affordable affordable and almost commonplace. Many applications have been developed such as teleconferencing, surgery, information assistance systems, games, and entertainment. To solve these different problems, different approaches for both animation control and modeling have been developed.

Active Part-Decomposition, Shape and Motion Estimation of Articulated Objects: A Physics-based Approach
Ioannis A. Kakadiaris
Dimitri Metaxas
Ruzena Bajcsy
MS-CIS-94-57
GRASP LAB 381

We present a novel, robust, integrated approach to segmentations shape and motion estimation of articulated objects. Initially, we assume the object consists of a single part, and we fit a deformable model to the given data using our physics-based framework. As the object attains new postures, we decide based on certain criteria if and when to replace the initial model with two models. These criteria are based on the model's state and the given data. We then fit the models to the data using a novel algorithm for assigning forces from the data to the two models, which allows partial overlap between them and determination of joint location. This approach is applied iteratively until all the object's moving parts are identified. Furthermore, we define new global deformations and we demonstrate our technique in a series of experiments, where Kalman filtering is employed to account for noise and occlusion.

Active Motion-Based Segmentation of Human Body Outlines
Ioannis A. Kakadiaris
Dimitri Metaxas
Ruzena Bajcsy
MS-CIS-94-58
GRASP LAB 382

We present an integrated approach towards the segmentation and shape estimation of human body outlines. Initially, we assume that the human body consists of a single part, and we fit a deformable model to the given data using our physics-based shape and motion estimation framework. As an actor attains different postures, new protrusions emerge on the outline. We model these changes in the shape using a new representation scheme consisting of a parametric composition of deformable models. This representation allows us to identify the underlying human parts that gradually become visible, by monitoring the evolution of shape and motion parameters of the composed models. Based on these parameters, their joint locations are identified. Our algorithm is applied iteratively over subsequent frames until all moving parts are identified. We demonstrate our technique in a series of experiments with very encouraging results.

Typing untyped λ-terms, or Reducibility strikes again!
Jean Gallier
MS-CIS-94-59
LOGIC & COMPUTATION 88

It was observed by Curry that when (untyped) λ-terms can be assigned types, for example, simple types, these terms have nice properties (for example, they are strongly normalizing). Coppo, Dezani, and Veneri, introduced type systems using conjunctive types, and showed that several important classes of (untyped) terms can be characterized according to the shape of the types that can be assigned to these terms. For example, the strongly normalizable terms, the normalizable terms, and the terms having head-normal forms, can be characterized in some systems $\mathcal{D}$ and $\mathcal{D}\otimes$. Our proofs use a new and more modular version of the reducibility method. As an application of our metatheorems, we show how the characterizations obtained by Coppo, Dezani, Veneri, and Pottinger, can be easily rederived. We also characterize the terms that have weak head-normal forms, which appears to be new. We conclude by stating a number of challenging open problems regarding possible generalizations of the realizability method.

Proving Properties of Typed λ-Terms Using Realizability, Covers, and Sheaves
Jean Gallier
MS-CIS-94-60
LOGIC & COMPUTATION 89

The main purpose of this paper is to take apart the reducibility method in order to understand how its pieces fit together, and in particular, to recast the conditions on candidates of reducibility as sheaf conditions. there has been a feeling among experts on this subject that it should be possible to present the reducibility method using more semantic means, and that a deeper understanding would then be gained. This paper gives mathematical substance to this feeling, by presenting a generalization of the reducibility method based on a semantic notion of realizability which uses the notion of a cover algebra (as in abstract sheaf theory). A key technical ingredient is the introduction a new class of semantic structures equipped with preorders, called pre-applicative structures. These structures need not be extensional. In this framework, a general realizability theorem can be shown. Kleene's recursive realizability and a variant of Kreisel's modified realizability both fit into this framework. We are then able to prove a meta-theorem which shows that if a property of realizers satisfies some simple conditions, then it holds for the semantic interpretations of all terms. Applying this theorem to the special case of the term model, yields a general theorem for proving properties of typed λ-terms, in particular, strong normalization and confluence. This approach clarifies the reducibility method by showing that the closure conditions on candidates of reducibility can be viewed as sheak conditions. the above approach is applied to the simply-typed λ-calculus (with types $\rightarrow, \times, +$, and $\perp$), and to the second-order (polymorphic λ-calculus (with types $\rightarrow$ and $\forall^2$), for which it yields a new theorem.

Linear Structure From Motion
Inigo Thomas
Eero Simoncelli
MS-CIS-94-61
GRASP LAB 383

Determining the structure of the world and the motion of the observer from image changes has been a central problem in computer vision for over fifteen years. Since the early work on *Structure from Motion* (SFM) by Longuet-Higgins[4] and Pradzny[6], several techniques have been developed to compute the motion of the camera, the shape of moving objects, or distances to points in the world. However, the image changes are *non-linearly* related to camera motion and distances to points in the

world. Thus, solving the problem typically requires non-linear optimization techniques that can be unstable or computationally inefficient. Linear algorithms are preferable since they are computationally advantageous, and since linear estimation is much better understood than non-linear estimation. Our paper describes an unbiased, completely linear algorithm for Structure-from-Motion. This work is similar to that of Jepson & Heeger [3] except that we employ spherical projection. The use of a spherical imaging geometry allows a simpler and more intuitive derivation of the algorithm, and produces an unbiased estimator. Experimental results are provided that demonstrate the performance of the algorithm.

Combining color and geometry for the active, visual recognition of shadows
Gareth Funka-Lea
Ruzena Bajcsy
MS-CIS-94-62
GRASP LAB 384

In computer vision for object recognition or navigation, shadows are a frequent occurrence. However, shadows are difficult to recognize because they cannot be infallibly recognized until a scene's geometry and lighting are known. We present a number of cues which *together* strongly suggest the identification of a shadow and which can be examined without a high computational cost. The techniques developed are: a color model for shadows and a color image segmentation method that recovers single material surfaces as single image regions irregardless of whether the surface is partially in shadow; a method to recover the penumbra and umbra of shadows; and a method for determining whether some object could be obstructing a light source. These cues either depend on or their reliability improves with the examination of some well understood shadows in a scene. Our observer is equipped with an extendable probe for casting its own shadows. These actively obtained shadows allow the observer to experimentally determine the number, location, and rough extent of the light sources in the scene. The system has been tested against a variety of indoor and outdoor environments.

Planning and Terrain Reasoning
Michael B. Moore
Christopher Geib
Barry D. Reich
MS-CIS-94-63
LINC LAB 280

We describe the ZAROFF system, a plan-based controller for the player who is "it" in a game of hide and seek. the system features visually realistic human figure animation including realistic human locomotion. We discuss the planner's interaction with a changing environment to which it has only limited perceptual access.

Human Body Simulations

FINAL REPORT

Norman Badler
Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389
215-898-5862
badler@central.cis.upenn.edu

1 Summary

The software system *Jack* for human figure modeling and animation is the product of many years of sponsored research, much of it from U.S. Army sources. This report summarizes the culmination of that effort by viewing the theoretical and algorithmic underpinnings of the research, the application domains and their impact, and recent steps towards commercialization of the technology.

2 Background

People are all around us. They inhabit our home, workplace, entertainment, and environment. Their presence and actions are noted or ignored, enjoyed or disdained, analyzed or prescribed. The very ubiquity of other people in our lives poses a tantalizing challenge to the computational modeler: people are at once the most common object of interest and yet the most structurally complex. Their everyday movements are amazingly fluid yet demanding to reproduce, with actions driven not just mechanically by muscles and bones but also cognitively by beliefs and intentions. Our motor systems manage to learn how to make us move without leaving us the burden or pleasure of knowing how we did it. Likewise we learn how to describe the actions and behaviors of others without consciously struggling with the processes of perception, recognition, and language.

Our research on human body modeling and simulation has been done in the Computer Graphics Research Lab of the Computer and Information Science Department at the University of Pennsylvania since approximately 1975. We have achieved international recognition for our research and specifically the *Jack*®software. Much of the Lab's work is reported in a recent book from Oxford University Press: *Simulating Humans: Computer Graphics, Animation, and Control*, by Badler, Phillips and Webber. In January 1994, we became the *Center for Human Modeling and Simulation* (HMS), a formal University entity with a Director (Norm Badler). Badler alone has produced 25 PhD students, and currently over 25 PhD students are active in the HMS Center's research. This HMS Center educates a large fraction of the PhD students in the CIS Department!

3 Theoretical Underpinnings

A famous Computer Scientist, Alan Turing, once proposed a test to determine if a computational agent is intelligent. In the Turing Test, a subject communicates with two agents, one human and one computer, through a keyboard which effectively restricts interaction to language. The subject attempts to determine which agent is which by posing questions to both of them and guessing their identities based on the "intelligence" of their answers. No physical manifestation or image of either agent is allowed as the process seeks to establish abstract "intellectual behavior," thinking, and reasoning. Although the Turing Test has stood as the basis for computational intelligence since 1963, it clearly omits any potential to evaluate physical actions, behavior, or appearance.

Later, Edward Feigenbaum proposed a generalized definition that included action: "Intelligent action is an act or decision that is goal-oriented, arrived at by an understandable chain of symbolic analysis and reasoning steps, and is one in which knowledge of the world informs and guides the reasoning." We can imagine an analogous "Turing Test" that would have the subject watching the behaviors of two agents, one human and one synthetic, while trying to determine at a better than chance level which is which. Human movement enjoys a universality and complexity that would definitely challenge an animated figure in this test: if a computer-synthesized figure looks, moves, and acts like a real person, are we going to believe that it is real? On the surface the question almost seems silly, since we would rather not allow ourselves to be fooled. In fact, however, the question is moot though the premises are slightly different: cartoon characters are hardly "real," yet we watch them and properly interpret their actions and motions in the evolving context of a story. Moreover, they are not "realistic" in the physical sense – no one expects to see a manifest Mickey Mouse walking down the street. Nor do cartoons even move like people – they squash and stretch and perform all sorts of actions that we would never want to do. But somehow our perceptions often make these characters *believable*: they appear to act in a goal-directed way because their human animators have imbued them with physical "intelligence" and behaviors that apparently cause them to chase enemies, bounce off walls, and talk to one another. Of course, these ends are achieved by the skillful weaving of a story into the crafted images of a character. Perhaps surprisingly, the mechanisms by which motion, behavior, and emotion are encoded into cartoons is *not* by building synthetic models of little creatures with muscles and nerves. The requisite animator skills do not come easily; even in the cartoon world refinements to the art and technique took much work, time, and study. Creating such movements automatically in response to real-time interactive queries posed by the subject in our hypothetical experiment does not make the problem any easier. Even Turing, however, admitted that the intelligence sought in his original test did not require the computational *process* of thinking to be identical to that of the human: the external manifestation in a plausible and reasonable answer was all that mattered.

So why are we willing to assimilate the truly artificial reality of cartoons – characters created and moved entirely unlike "real" people – yet be skeptical of more human-like forms? This question holds the key to our physical Turing Test: as the appearance of a character becomes more human, our perceptual apparatus demands motion qualities and behaviors which sympathize with our expectations. As a cartoon character takes on a human form, the only currently viable method for accurate motion is the recording of a real actor and the tracing or transfer ("rotoscoping") of that motion into the animation. Needless to say, this is not particularly satisfying to the modeler: the motion and actor must exist prior to the synthesized result. Even if we recorded thousands of individual motions and retrieved them through some kind of indexed video, we would still lack the freshness, variability, and adaptability of humans to live, work, and play in an infinite variety of settings.

If synthetic human motion is to be produced without the benefit of prior "real" execution and still have a shot at passing the physical Turing Test, then models must carefully balance structure, shape, and motion in a compatible package. If the models are highly simplified or stylized, cartoons or caricatures will be the dominant perception; if they look like humans, then they will be expected to behave like them. How to accomplish this without a real actor showing

the way is the challenge we address.

Present technology can approach human appearance and motion through computer graphics modeling and three-dimensional animation, but there is considerable distance to go before purely synthesized figures trick our senses. A number of promising research routes can be explored and many are taking us a considerable way toward that ultimate goal. By properly delimiting the scope and application of human models, we can move forward, not to replace humans, but to substitute adequate computational surrogates in various situations otherwise unsafe, impossible, or too expensive for the real thing.

The goals we set for ourselves are realistic but no less ambitious than the physical Turing Test: we seek to build computational models of human-like figures which, though they may not trick our senses into believing they are alive, nonetheless manifest animacy and convincing behavior. Towards this end, we

- Create an interactive computer graphics human model.
- Endow it with reasonable biomechanical properties.
- Provide it with "human-like" behaviors.
- Use this simulated figure as an agent to effect changes in its world.
- Describe and guide its tasks through natural language instructions.

There are presently no perfect solutions to any of these problems, but significant advances have enabled the consideration of the suite of goals under uniform and consistent assumptions. Ultimately, we should be able to give our surrogate human directions that, in conjunction with suitable symbolic reasoning processes, make it appear to behave in a natural, appropriate, and intelligent fashion. Compromises will be essential, due to limits in computation, throughput of display hardware, and demands of real-time interaction, but our algorithms aim to balance the physical device constraints with carefully crafted models, general solutions, and thoughtful organization.

From modeling realistic or at least reasonable body size and shape, through the control of the highly redundant body skeleton, to the simulation of plausible motions, human figures offer numerous computational problems and constraints. Building software for human modeling and simulation serves a widespread user population. In fact, it appears that such software has broad application since the features needed for analytic applications – such as multiple simultaneous constraints – provide extremely powerful features for a variety of users with contemporary computer skills but not necessarily animation talents. Our software design has tried to take into account a wide variety of physical problem-oriented tasks, rather than just offer a computer graphics and animation tool for the already skilled or computer-sophisticated animator.

Our research has built a solid underpinning of general techniques and theory for these ends:

- User interaction with 3D environments and especially articulated figures. The *Jack* user interface is one of its great attractions.

- Inverse kinematics running in real-time provides the underlying engine that permits manipulation of spatial constraints of a complex articulated figure, rather than requiring the user to manipulate individual joints directly.

- Reliance on open data databases for human factors information. A user can, if desired, customize human capability, size, and characteristics to an application. The Army ANSUR data is the default human anthropometric database. The joint limit and strength data are mostly from NASA sources.

- Human behaviors are often goal-directed, and sets of parallel, interacting goals are the normal way to control movement.

- Balance maintenance triggers a numbers of natural human behaviors such as counterbalancing, stepping, turning, and gait modification.

- Much of the realism in human motion comes from strength load and distribution considerations. We have successfully demonstrated natural synthetic motions based on strength and load for lifting as well as locomotion tasks.

- Simulated sensors are needed to achieve realistic interaction with an environment. *Jack* includes self-collision avoidance as well as real-time environmental object avoidance.

- Human gesture and facial expressions are not arbitrary, but are shaped by the agent's communicative goals and intentions. We have linked an underlying planning and question-answering system with realistic human gestures and facial expressions as the queries and responses are uttered by a speech generation system. These mappings are based on cognitive science principles.

- Synthesizing known or novel motions in real-time is a prerequisite for embedding realistic human agents into Virtual Environments or Distributed Interactive Simulations. *Jack* is the first convincing demonstration of this capability in both the Army ISMS and the Navy TTES systems.

- Natural Language instructions have been largely ignored by the AI community as a research subject; we have changed that by considering the role, requirements, and impact of NL processing of instructions as animation specifications.

The *Jack* testbed provides a highly capable vehicle with which to carry out research in 3D human-computer interaction, human animation, and especially the link between language and action. An important goal of our work is to enable people to use Natural Language instructions to guide the behavior of semi-autonomous agents. These are agents with behaviors and skills of their own (from low-level behaviors such as walking, obstacle avoidance, and terrain navigation, to higher-level skills such as the ability to manipulate objects for particular ends), that a person could enlist in carrying out a task. The situation is analogous to that of the leader of a group and its team members: the leader may issue instructions and specify policies affecting how those instructions are to be carried out. The team members should carry out the orders in

a situationally-sensitive manner that is consistent with policy as well as their own low-level behavioral responses. The instructions themselves should reflect the leader's awareness of the team members' skills and behaviors. By linking Natural Language to what is essentially real activity in a real world, we are able to experiment with theories concerning the situated interpretation of language in ways that theoreticians do not have access to.

4 Impact

We perform basic and applied research in human modeling and simulation. A number of particularly well-motivated application areas are:

- Human factors analysis.
- Physical performance assessment.
- Intelligent agents acting on instructions.
- Maintenance/logistics technician.
- Factory or assembly worker.
- Dismounted, semi-autonomous soldier.
- Virtual training simulator.
- Medic or surgical assistant.
- Virtual human for surgical procedures,

While not as exciting as motion picture characters, as personable as cartoons, or as skilled as Olympic athletes, there are justifiable uses to virtual human figures in these domains. Visualizing the appearance, capabilities and performance of humans is an important and demanding application. The lessons learned may be transferred to other problems, applications, and even more entertaining uses of human-like models.

Besides *Jack*'s immediate applicability to directing virtual agents in virtual environments, we can begin to create animated simulations from information-efficient task specifications. Bonnie Webber's research will make significant contributions to theoretical and computational linguistics by studying:

- How *instruments* are specified in Natural Language instructions and how agents recognize from these instructions and/or the world, ways in which instruments will support their intended activity;
- *Policy specification* and understanding, aimed at developing a framework for planning and sensing that enables policy to inform and/or constrain behavioral choice; and
- How instructions specify task-appropriate *sensing behavior* – what sensing modalities the agent should employ, what is being sensed for, and how often to perform sensing, since successful perception may require other physical activities to be performed as well.

5 Commercialization

Since 1983 we have been required to and actually have delivered software to our sponsors. The first interstate delivery of the *Jack* software was made to NASA Johnson Space Center in 1987. The U.S. Army Human Engineering Lab at Aberdeen has been a major sponsor, user, and partner in the development of this software. In early 1990 the University of Pennsylvania initiated a formal trademark request for the name *Jack* as a software product. Recently, trademark protection has been extended to the EC. The actual commercial distribution of the *Jack* software began through the University's Center for Technology Transfer (CTT) in 1992. With the creation of the HMS Center, all contacts are now funneled through our office with the exception of legal documents which are generated or reviewed by CTT and/or the Office of Research Administration.

One portion of *Jack* has been embedded into other commercial software. *Jack*'s inverse kinematics algorithm is the backbone of Wavefront Technologies *Kinemation* product. *Jack* enjoys a direct interface with Micro Analysis and Design's *HOS* software and Technomatix's CAD system. We have built CAD geometry translators (most under Army contract) into and out of *Jack* for most major CAD systems. These are supplied at no additional cost to any *Jack* system requiring them. Other interface or embedding agreements are underway with several major CAD vendors and other software interests.

Modeling Clothed Figures
Final Report: Army Grant DAAL03-90-G-0191

Norman I. Badler
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389

August 31, 1992

Clothing Models

Abstract

In most workplace environments we have encountered, clothed figures are the norm and would be expected by the designer. Adding clothing to a human figure improves its graphical appearance and realism. Clothes modeling can be done in many ways ranging from very simple to more realistic but complicated. The simplest technique is to change the attributes of certain segments of the body figure; for example, by modifying the colors of the lower legs we get the effect of a body wearing short pants. This is not quite as silly as it sounds, because the body segment geometry can be created with a clothed rather than bare-skinned shape. The best but more complicated approach is to drape and attach clothing over a body to simulate the intricate properties of garments.

Besides improving realism, there is a practical human factors aspect to clothing. Clothing constrains movement by restricting the joint angle limits. An approach to analyzing this problem has been developed by using collision detection for a geometric clothes model.

1 Geometric Modeling of Clothes

Rigid models for clothing are created by specially designing psurfs on a segment-by-segment basis. Thus a shirt or jacket would have, say, five parts: one for the torso, and two for each limb segment. Clothing exists independently of a given figure model as a library of objects which can be selectively placed on a model at user determined sites. This data base is modifiable through typical geometric editing commands. A clothing item is positioned by matching key points on the clothing to key points on the body segments. A global deformation algorithm can be used to fit the clothing piece correctly on the segment.

One apparent problem with geometrically modeled clothing occurs when the human figure moves joints. Since the clothing model is not deformable, there are gaps between segments. (This is in fact true even without clothing if the figure is modeled with polyhedral meshes. As the geometry is carried on the segment, it inherits the geometric transformation without any compensation for the interaction of material, flesh or clothes, at the joint.) Extra work is necessary to alleviate the joint gap problem. A gap filling algorithm has been developed to make up these gaps when animating. It connects the boundaries of two adjacent segments by generating spline surfaces using the tangent information of the two segments at the boundaries.

As an initial attempt to develop geometric models of clothes, a sweatshirt and pants were designed (Figure 1). Each segment of a clothing item is a psurf whose geometry and position are closely related to a corresponding body segment. The following is a step by step procedure for geometric clothes design.

1. Determine the size of clothes:

   In conventional clothing design, circumferences are measured at certain positions of the body in order to determine the clothing size. In our approach the maximum segment breadth in two orthogonal directions are measured instead of circumferences.

   The following list shows the positions at which the maximum breadths are measured and lists their corresponding slices from the accurate biostereometric slice bodies.

   - For trousers

waist	the 6th slice of lower_torso
hip	the 9th slice of hip_flap
upper leg	the first slice of upper_leg

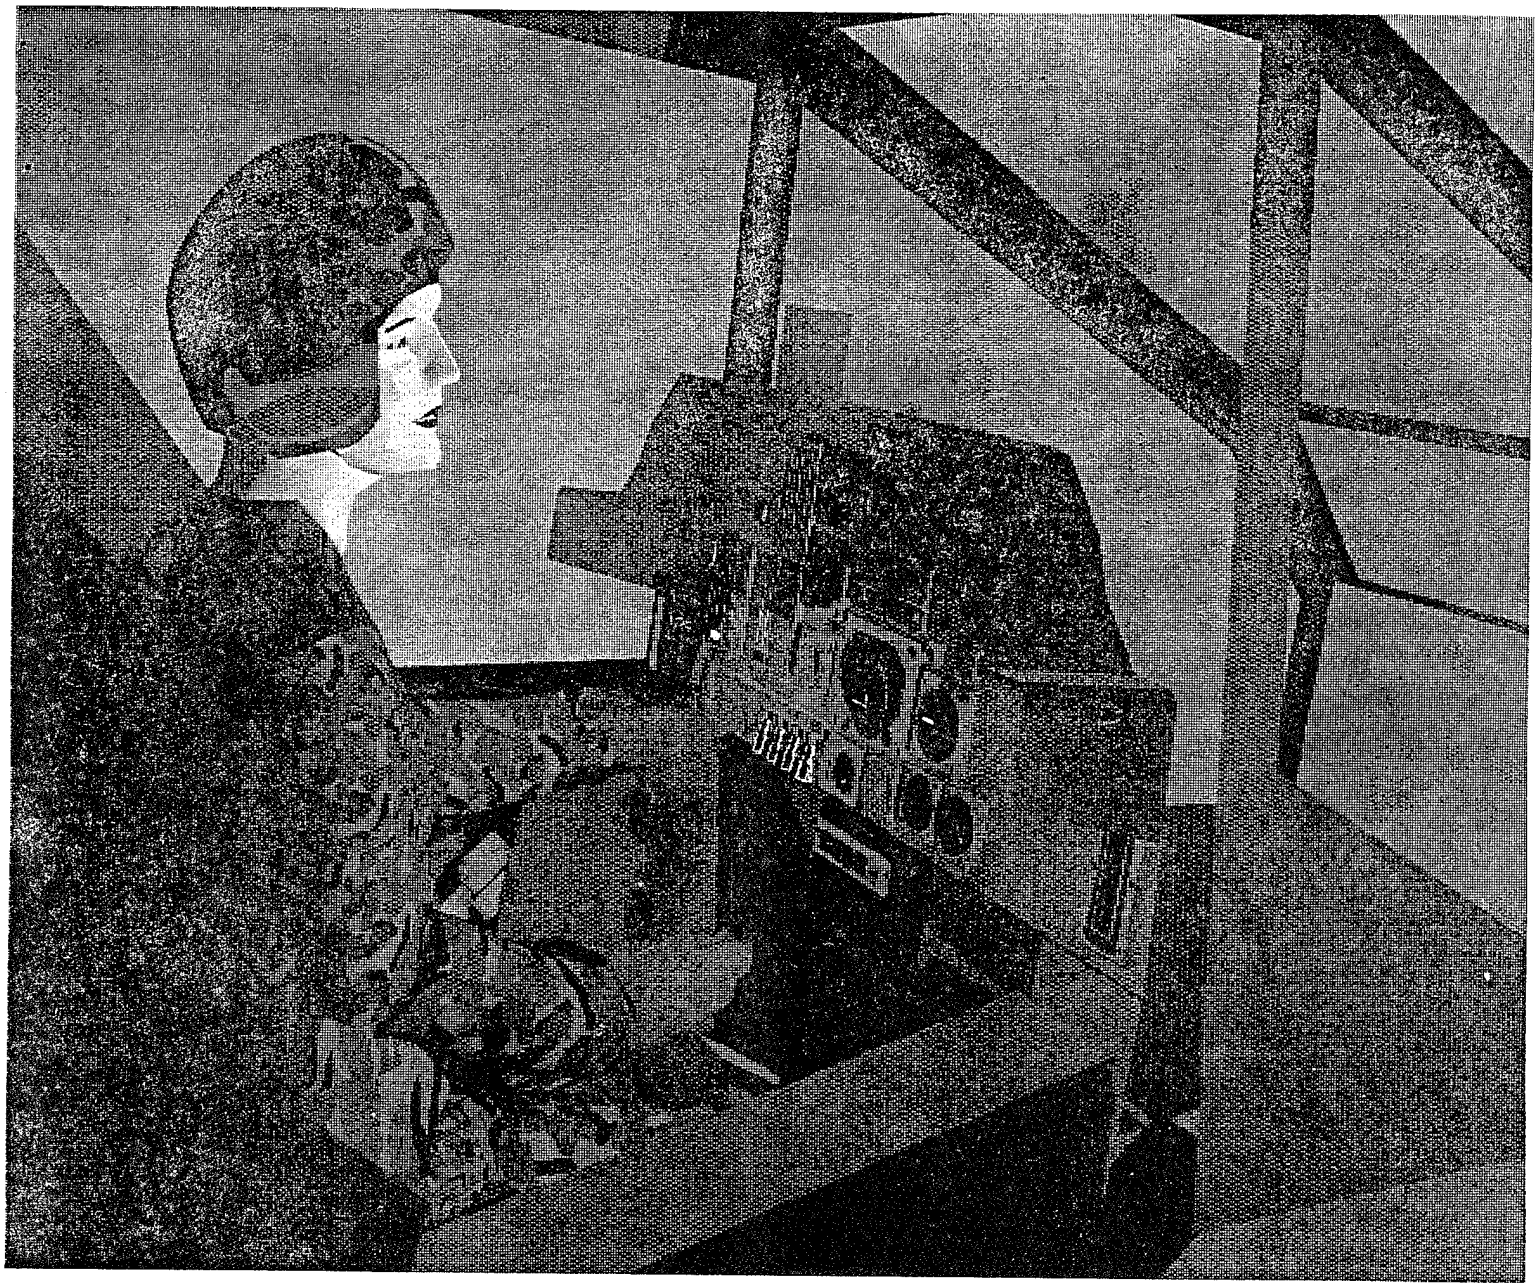

Figure 1: Texture-Mapped Camouflage Clothing on Contour Body Figure in Apache Helicopter Cockpit. Cockpit Courtesy NASA Ames Research Center. Image by Eunyoung Koh, Pei-Hwa Ho, and Jiahe Lu.

lower leg	the last (24th) slice of lower_leg

- For a sweatshirt

neck	the first slice of upper_torso
chest	the 6th slice of upper_torso
breast	the 12th slice of upper_torso
waist	the first slice of lower_torso
upper arm	the 8th slice of upper_arm
lower arm	the last (16th) slice of lower_arm

Also, a measure of the length from the neck to the bottom of the shirt is provided to determine the length of the shirt. These sizes are specified in advance for each article of clothing.

2. Create psurfs:

   The geometry of each segment is determined by a set of body slices. A new set of slices for clothing is constructed by sculpturing the body slices depending on the shape of the segment and the specified sizes in the previous step. The fundamental idea in the construction is to pick a few thick slices and duplicate them appropriately along the segment after scaling. Scaling is done by linear interpolation so that the scaled slices may match with the specified maximum breadth sizes at the positions designated.

   The completed surface definition of a clothes segment can be obtained by tiling the slices. Tiling is performed by generating rectangles from the data points which define two adjacent slices.

3. Attach clothes segments to human body:

   Each clothes segment can be attached to the corresponding body segment by a joint which is located at the upper part of that segment.

The clothing shape can be easily modified by changing the slice definition of the clothes. For example, folded sleeves, short sleeves, and short pants can be simulated by simple modification or deletion of slices.

2 Draping model

The most realistic clothing can be created by simulating the support and draping of pattern pieces of arbitrary shape. Wrinkling, folding, and the effects of gravity are displayed through a relaxation method or a finite element method. Pattern pieces may also be stitched at seams and draped simultaneously. Pattern pieces of different lengths may be sewn together, resulting in an oversewing effect.

The draping of the pattern pieces is done on a figure in a static posture. Interference testing is done in the draping algorithm to make sure that the pattern pieces slide over the surface of the figure without penetrating the surface (Figures 2).

There are several methods to simulate the draping of a square piece of cloth, isolated from other cloth, which are based on a relaxation method. Feynman [Fey86] uses a formula which minimizes the energy of a cloth and tries to simulate the shape of thin flexible membranes under the influence of force fields and rigid bodies. The local minimum of the cloth is found by moving each of the grid points in turn toward a position which decreases the energy of the cloth. The energy expression of a cloth is described as:

$$E_{total}(S) = k_s s(S) - k_b b(S) - k_g g(S)$$

where $s(S), b(S), g(S)$ represent the effects of strain, bending, and gravity. The parameters k_s, k_b, k_g control the relative strengths of these three effects: a large k_s means the cloth is difficult to stretch; a large k_b means the cloth is stiff and resists bending; and a large k_g means the cloth is heavy.

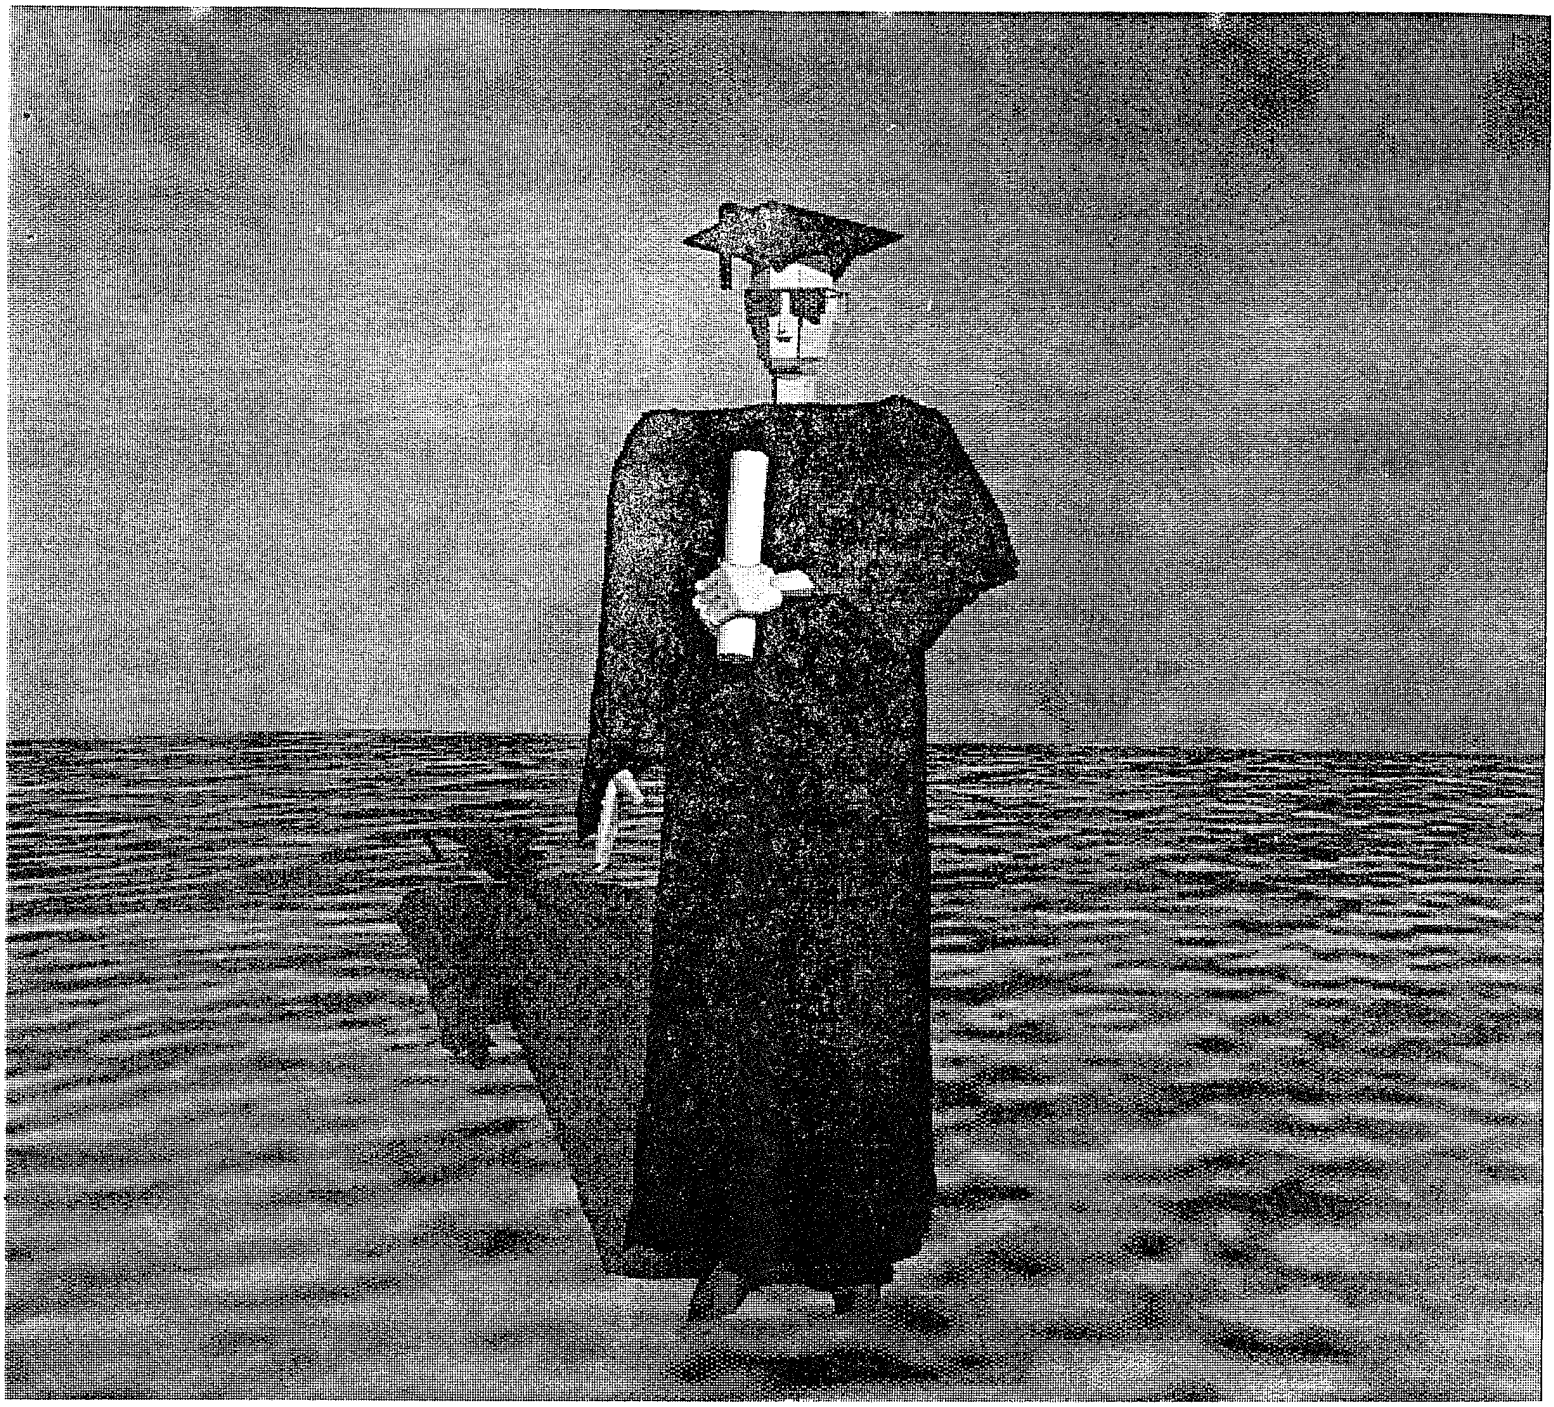

Figure 2: Graduation Day. Cloth modeling by Lauren Bello. Draping model by Welton Becket.

Relaxing a single point is the process of moving it so that the energy of the cloth of which it is a part is decreased. The method used to relax a single point first finds the direction in which the point would most like to move: the direction of the negative gradient of the energy as a function of position. Then it moves the single point in that direction so that its energy is minimized.

Feynman suggests using a multigrid method to speed up the relaxation sweeping process. However, it must be used carefully to avoid distortion. He also introduces fixed points in order to forbid the cloth to move into a solid.

Weil [Wei86] considered the problem of hanging a piece of cloth by fixing some constraint locations of the cloth. The cloth is represented as a rectangular grid (u, v) of 3D coordinates (x, y, z). His method is a two phase algorithm. The first part approximates the surface within the convex hull in (u, v) space of the constraint points; that is, all the interior points are placed on catenaries. The second phase uses an iterative relaxation process to minimize maximum displacement of all the points in the grid up to a given tolerance.

The approximation process first positions the constraint points, then traces catenaries between each pair of constraint points, thereby positioning all grid points in such a catenary and creating triangles. If a point is in the path of two catenaries, it selects the highest point because it is a constraint that can't be violated (catenaries are the lowest possible points). All triangles created are subdivided by tracing catenaries from each vertex of the triangle to the midpoint of the opposite edge and creating two new triangles, that are recursively subdivided until all interior points have been positioned.

The relaxation phase tries to propagate the displacement of grid points until all maximum displacements fall below a certain tolerance. Determining the displacement of a point is done by obtaining the displacement vectors to position the point at the correct distance from each of its neighbors and adding the vectors. The final displacement is the square root of the sum of the squares of the displacements so that larger displacements will have more influence.

Terzopoulos, Platt, Barr and Fleisher [TPBF87] use elasticity theory to describe the behavior of a deformable object. The model responds in a natural way to applied forces, constraints, and impenetrable obstacles. The equations of motion governing the dynamics of the deformable bodies under the influence of applied forces is given

$$\frac{\partial}{\partial t}(\mu \frac{\partial r}{\partial t}) + \gamma \frac{\partial r}{\partial t} + \frac{\delta \mathcal{E}(\nabla)}{\delta \nabla} = f(r, t),$$

where $r(a, t)$ is the position of the particle a at time t, $\mu(a)$ is the mass density of the body at a, $\gamma(a)$ is the damping density, and $f(r, t)$ represents the net externally applied forces. $\mathcal{E}(\nabla)$ is a functional which measures the net instantaneous potential energy of the elastic deformation of the body.

To create animation with this model, the motion equation is solved numerically, integrating through time. This model is active in the sense that it responds to forces and interacts with objects.

3 Computing Joint Limits

We tried to test the impact of clothing items on joint limits. A new joint limit affected by a clothing item is estimated by detecting collision among segments. Note that two adjacent segments will always give a collision. In order to solve this problem, we introduced a collision distance threshold: any nodes within the threshold from the joint position are ignored during collision detection.

We experimented with geometric models of an army vest and armor plate, by bending the left arm (left elbow joint) against the chest under three conditions: with the vest on, with armor plate attached, and with no clothes (Figure 3 and 4). The computed elbow joint limits were 78.5 deg, 96.0 deg, and 159 deg, respectively. Similar experiments with shirts of different thicknesses showed different effective joint limits of the left elbow joint when the figure tried to bend its arm.

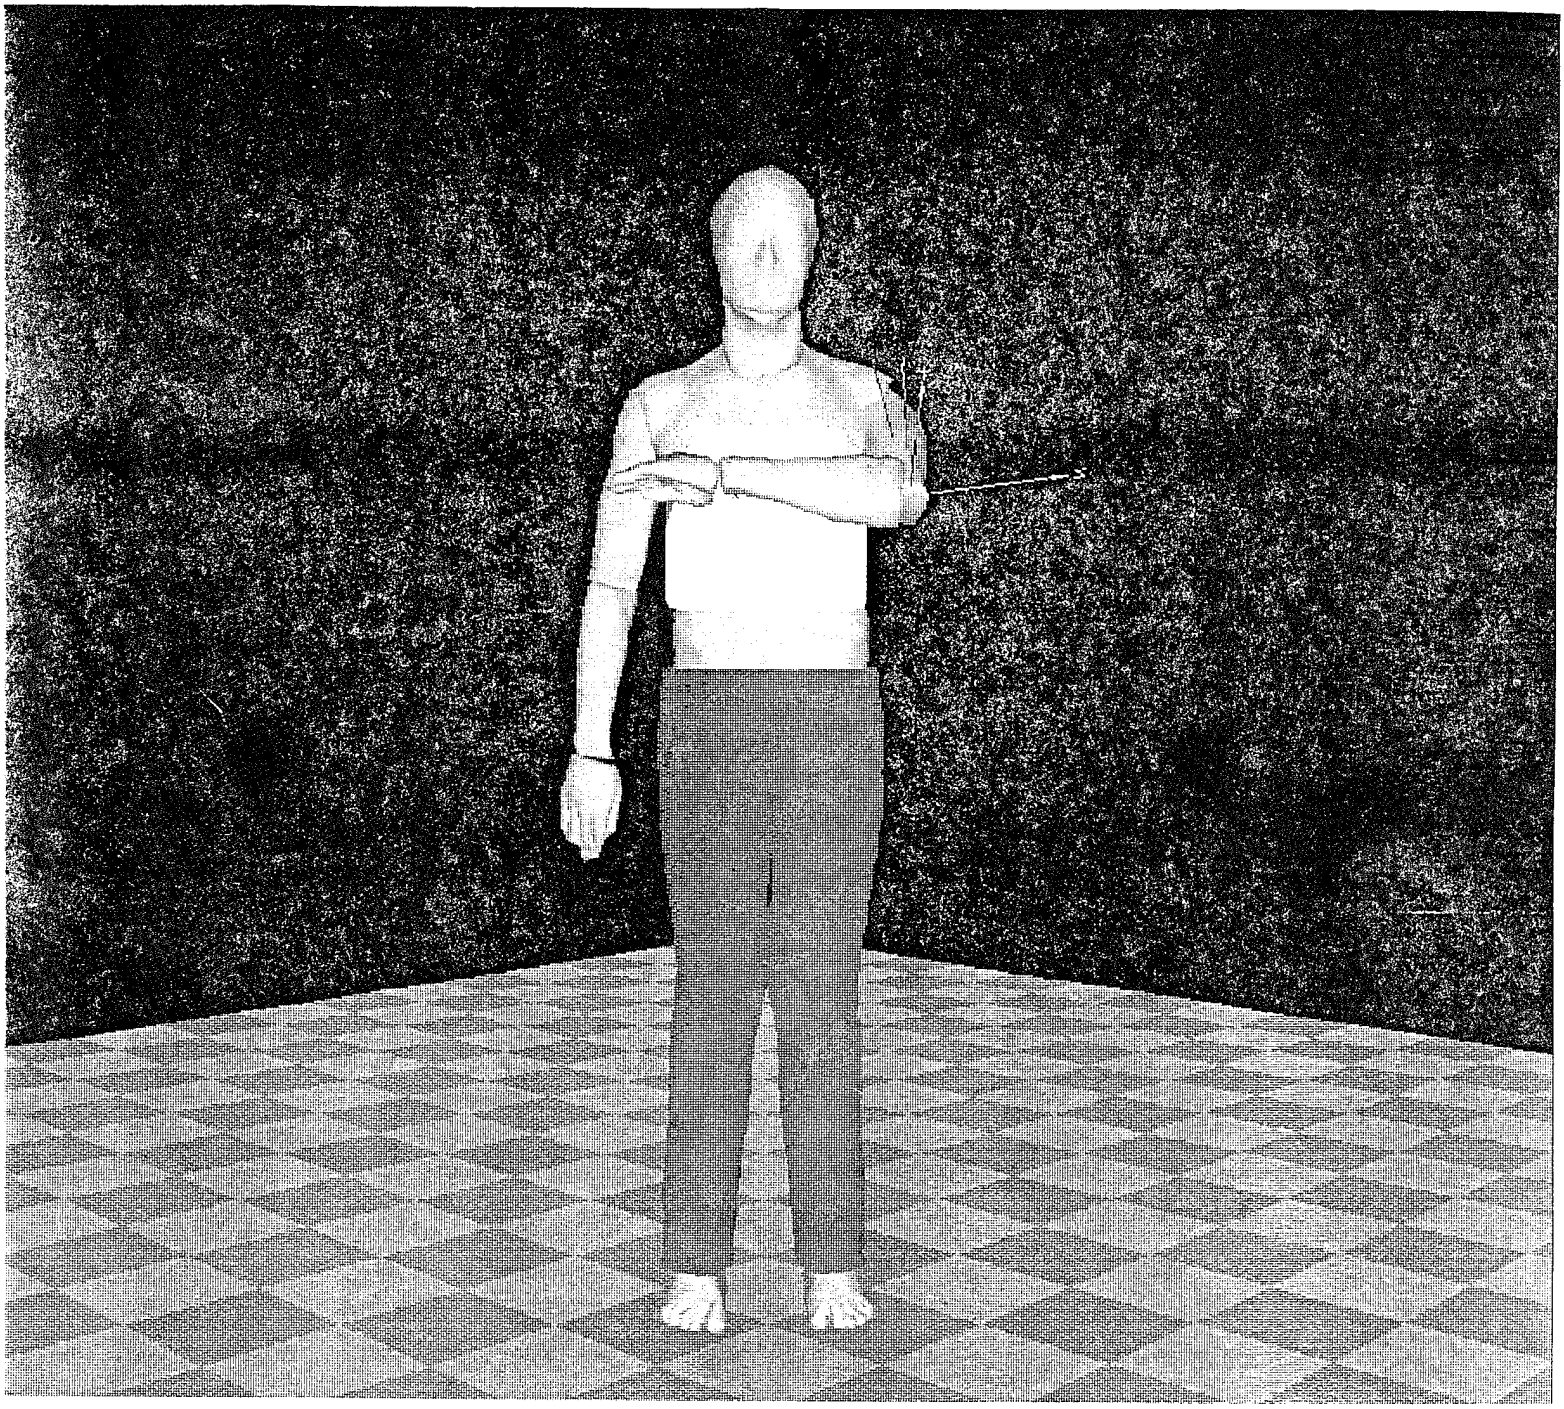

Figure 3: Army armor plate clothes for mobility experiments. Modeling by Eunyoung Koh.

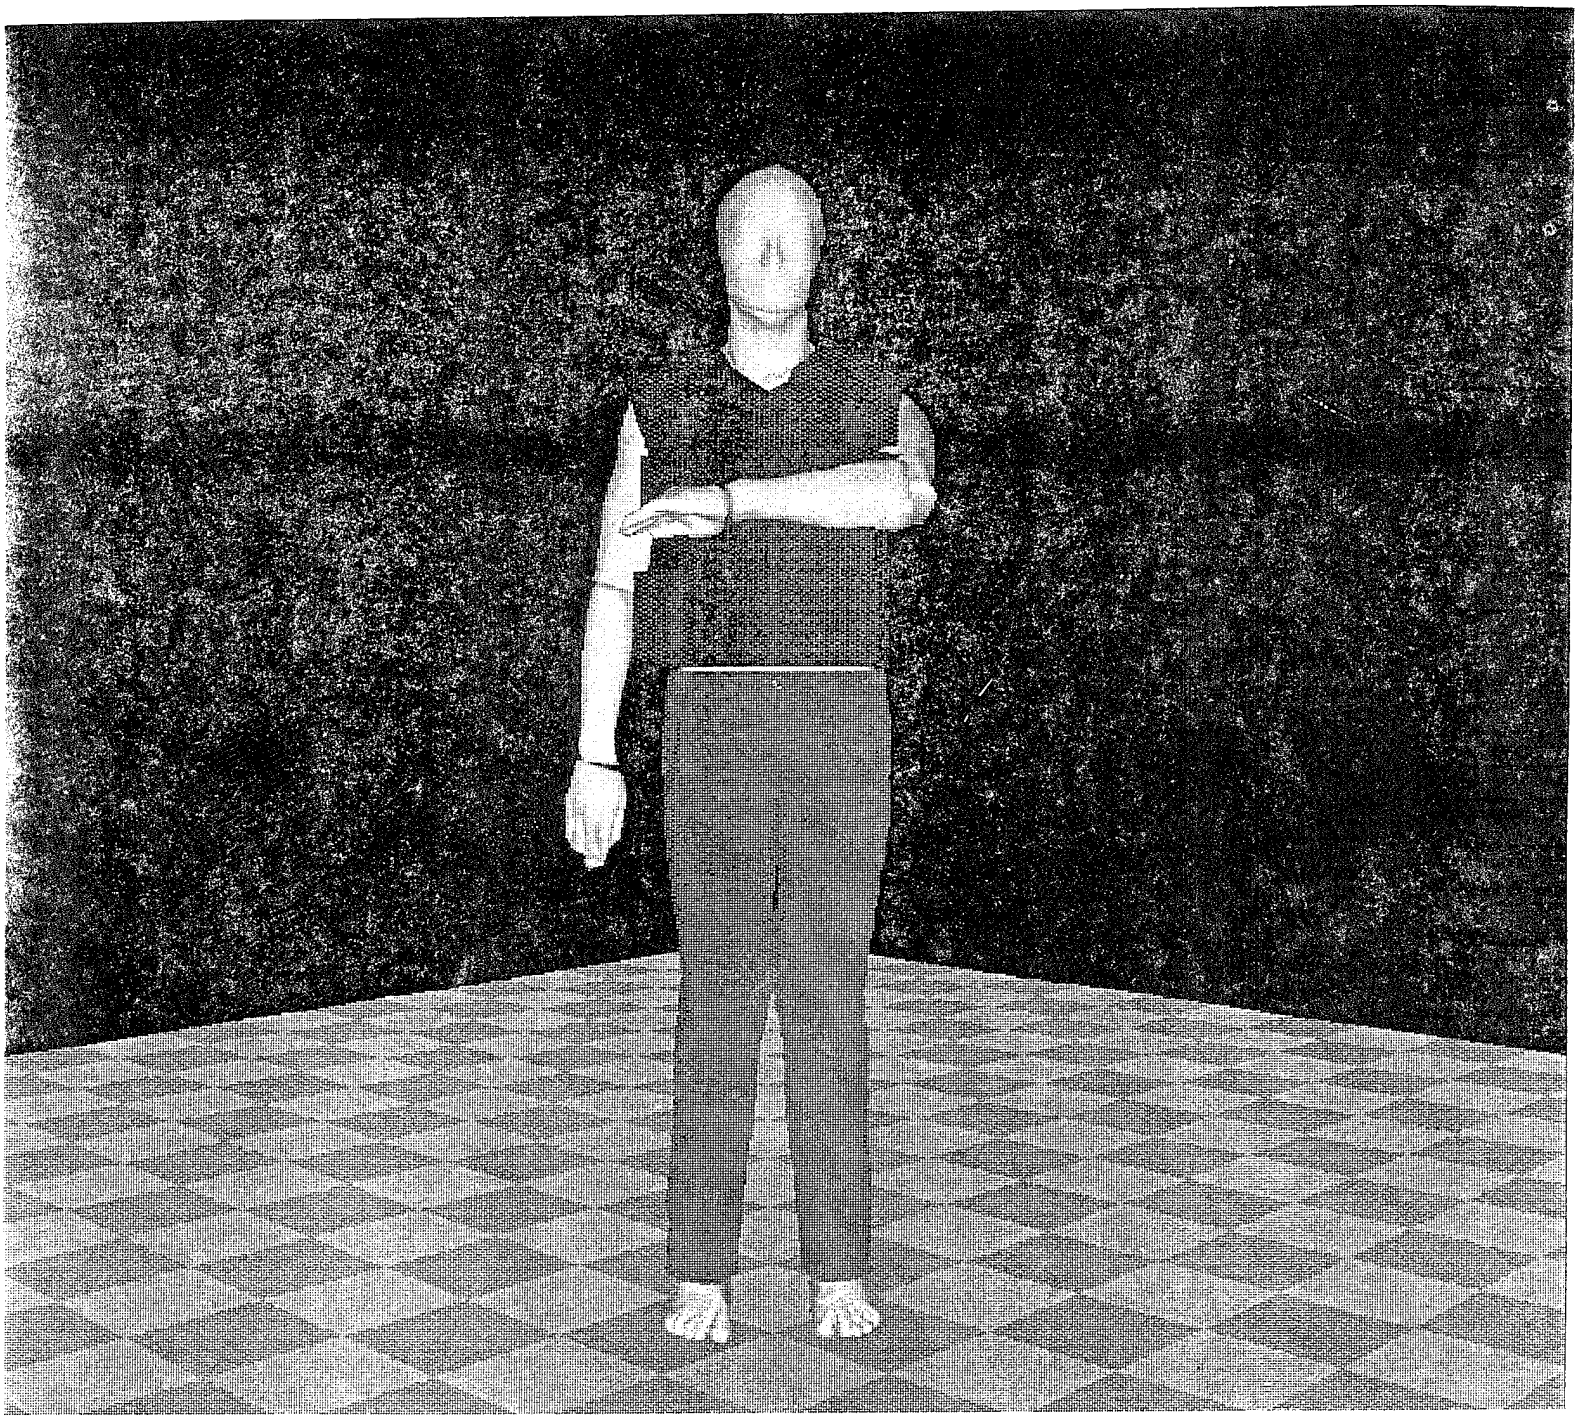

Figure 4: Army vest (with armor underneath) for mobility experiments. Modeling by Eunyoung Koh.

4 References

References

[Fey86] Carl R. Feynman. Modeling the appearance of cloth. Master's thesis, Massachusetts Institute of Technology, 1986.

[TPBF87] Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. Elastically deformable models. In *Computer Graphics*, volume 21, pages 205–214, 1987.

[Wei86] Jerry Weil. The synthesis of cloth objects. *Computer Graphics*, 20(4):49–54, 1986.

Human Factors Simulation Research — In Cooperation with NASA Ames Aeroflight Directorate A^3I Project

FINAL REPORT

Norman Badler
Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389
215-898-5862
badler@central.cis.upenn.edu

1 Summary

Enhancements to the *Jack* model and software for the A^3I effort focused on three areas: coupled eye and head motion, better grasping actions, and improvements in collision detection.

2 Results

1. We added visible eyes with optional eyeglasses to the *Jack* figure. We extended the dependent joint capability in *Jack* to couple head and eye movements. The general problem is complicated by the fact that there is no fixed relationship between eye gaze direction and head movement; rather, there is considerable hysteresis in the process. The eyes can move small angles without causing head motion, but when they exceed a threshold the head rotates to a new orientation. We constructed the eye-head motion dependency within *Jack* based on Sparks' [4] empirical data and model. Video samples of this eye and head motion are contained on the SAE Human Modeling Video [1].

2. The human figure's hands were articulated and visually improved. Preliminary grasping behaviors were added based on Iberall's taxonomy of handshapes [3, 2]. The hand functionality includes definitions for the major Iberall grasp types, grasping that continues while the grasped object is moved, and a "compliant grasp" that allows a sliding contact with collision-detection mediated joint angle adjustments during hand or object motion. A technique using point-to-point distances as parallel proximity "sensors" was used to do efficient collision detection on the hand and fingers [5].

3. In addition to hand grasps, whole body collision detection was provided. A novel feature of the implemented technique was that detected collisions created constraints that were then just added into the current behavioral constraint set for the figure. In this way the figure's collision response was smoothly integrated with the other *Jack* behaviors such as reach and balance.

References

[1] Norman I. Badler. "Jack: The Human Factor" on SAE "Human Modeling Technology and Applications" Video SP-1100, 1995.

[2] Norman I. Badler, Cary B. Phillips and Bonnie L. Webber. *Simulating Humans: Computer Graphics, Animation and Control.* New York: Oxford University Press, 1993.

[3] T. Iberall. The Nature of Human Prehension: Three Dextrous Hands in One. *IEEE Conference on Robotics and Automation*, 1987, pp. 396-401.

[4] David L. Sparks. "Neural control of orienting eye and head movements," Proceedings of the Dahlen Conference on Motor Control, Berlin, 1989.

[5] Xinmin Zhao and Norman I. Badler. "Near Real-Time Body Awareness," *Computer-Aided Design* 26(12), pp. 861-868, Dec. 1994

Human Modeling for the U.S. Army Tank and Automotive Command

FINAL REPORT

Norman Badler
Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389
215-898-5862
badler@central.cis.upenn.edu

1 Summary

We supplied U.S. Army TACOM with object code versions of the *Jack* software system for human figure modeling and animation. The principal activity undertaken with this grant is the integration of *Jack* into the CAD environment at TACOM.

2 Results

We determined the most appropriate way to convert BRL-CAD geometry files to the *Jack* surface models. Both TACOM and UPenn had partial BRL-CAD to surface geometry converters; the TACOM version ("Fred") was more robust, while the UPenn version permitted object articulation. A combination scheme was constructed and reported on at the 1991 BRL-CAD Symposium [1].

In addition to the static BRL-CAD geometry conversion, we attempted to construct a geometry and articulation converter from *Jack* into *DADS*, a dynamics simulation system used extensively at TACOM. We were able to produce the required conversions but the *DADS* software ran far too slowly on the highly articulated *Jack* human figure to be of any practical use. Subsequently we began implementation of our own dynamics simulation.

References

[1] Osman Niazi and Norman Badler. "Converting BRL-CAD objects to surface representation and adding articulation and *Jack* ergonomic analysis," BRL-CAD Symposium '91, May 1991.

New Techniques for Merging Database Schemas
Final Report - Subcontract to DAAL03-89-C-0031

Peter Buneman
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389
(215) 898-7703
peter@central.cis.upenn.edu

1 Results

The problem of merging database schemas arises from two tasks: *database integration*, a problem which is endemic to any large organization that supports multiple heterogeneous databases, and *integration of user views*, which occurs during the design phase of a database when constructing a schema that satisfies the individual needs of each of a set of user groups. Under this funding, we noted that these two kinds of schema integration are fundamentally different. The reason for this can be seen by considering the direction in which data is transformed in each case. For database integration, instances of each of the source databases are transformed into instances of the merged schema. On the other hand, when integrating multiple user views instances of the merged schema must be transformed back into instances of the user views. The intuition is that when integrating user views *all* of the underlying information must be represented; no objects or attributes can be missing since some user may want the information. However, when integrating pre-existing databases the best that can be hoped for is that attributes of objects that are present in every underlying database will be definitely be present in the integration; attributes that are present in some but not all of the underlying databases may be absent in the integration. In [1,3], it was observed that integrating user views corresponds to the "least upper bound" of the component schemas in some information ordering on schemas, while in database integration what is required is the "greatest lower bound" of the component schemas in some information ordering on schemas. A good schema-integration method should therefore take account of its intended purpose and include a semantics for the underlying transformations of instances.

Having developed a semantics for the type of database merging involved, we were able to construct a technique for merging database schemas that is simple, general, and does not suffer from "arbitrary" results that are endemic to existing techniques. The technique has been implemented with a GUI environment, as discussed in [2]. A shortcoming of this technique when applied to database integration, however, is that it does not capture the initial transformations that may be necessary to bring schemas to a form in which they can be merged. That is, similar concepts may be represented in dramatically different ways depending on how they are used in different databases and hence must be modified prior to merging. Furthermore, we have found that there is considerable interplay between database transformations and constraints. We therefore developed in [5] a declarative language based on Horn clause logic for expressing database transformations and constraints. This practical environment has proven enormously effective in applications associated with the Human Genome Project (see [4]).

2 References

Publications resulting from this grant:

1. "Theoretical Aspects of Schema Merging" by P. Buneman, S. Davidson and A. Kosky. *Extending Database Technology* (March 23-27, 1992), Springer-Verlag, pp. 152-167.

2. "A Basis for Interactive Schema Merging" by P. Buneman, S. Davidson and A. Kosky. *Proceedings of the 25'th Hawaii International Conference on System Sciences* (Jan. 7-10, 1992), pp. 311-322.

3. "Modeling and Merging Database Schemas" by A. Kosky. Technical Report MS-CIS-91-65/L&C 39, Department of Computer and Information Science, University of Pennsylvania, 1991.

4. "Facilitating Transformations in a Human Genome Project Database," with Anthony Kosky and Barbara Eckman. *3rd International Conference on Information and Knowledge Management* (November 1994).

5. "Types with Extents: On transforming and querying self-referential data-structures" by A. Kosky. PhD Thesis Proposal, Technical Report MS-CIS-95-21, University of Pennsylvania, May 1995.

Supporting the Dissemination of TraumAID: A Decision Aid for the Management of Injuries

Final Report - Subcontract to DAAL03–89–C–0031

Bonnie L. Webber, Ph.D.
Department of Computer and Information Science
University of Pennsylvania
Philadelphia PA 19104-6389
215-898-7745
bonnie@central.cis.upenn.edu

1 Summary

The aim of this work was to make TraumAID, the decision aid we had developed for the management of multiple trauma, more accessible both to a wider community of AI researchers and to medical personnel at sites beyond the Medical College of Pennsylvania. At the time of the proposal, TraumAID could only run on a Symbolics Lisp Machine, which was at the time, a common, but very expensive workstation. We proposed to develop a version of TraumAID that would be independent of its Symbolics Lisp Machine embodiment. This would involve two things:

1. isolating its user-interface from its purely internal computations (i.e., those of the reasoner and planner) and define the information and control flow between the two. The information and control flow should be compatible with the HyperCard interface we have been developing for the Mac II over the past two years. It should also be compatible with a purely functional X-Windows interface on the Sun (or other UNIX machine).

2. eliminating Symbolics Lisp-specific language features in the reasoner and planner, in favor of features that comply with the Common Lisp standard, possibly augmented with 'standard' extensions such as *flavors*.

We proposed to take both these steps during the period of the grant, demonstrating success, by having TraumAID running on one of the two low-cost workstations – either a Sun SparcStation 1 IPC or a Mac-IIfx.

2 Results

We successfully ported the core of TraumAID to both a SUN environment (in Lucid Common Lisp) and a Macintosh environment (in Mac Common Lisp). To do this, and to allow for additional future ports to other environments such as a PC or an SGI, we separated TraumAID's interface software from its reasoning/planning software and knowledge bases. On the SUN. TraumAID has a menu interface implemented in X-Windows, while for the Mac, two interfaces were developed – a simple menu-based interface, which we call "MacTraumAID" and a HyperCard interface, which we call "HyperTraumAID". All three interfaces have been integrated with TraumAID's reasoning/planning software and knowledge interfaces.

Both the Mac and Sun versions of TraumAID continue to be used, the Sun version for system development and the Mac version for clinical testing.

For additional information, TraumAID maintains a homepage on the World Wide Web at `http://www.cis.upenn.edu/ traumaid`.

Virtual Reality Explorations

FINAL REPORT

Norman Badler
Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389
215-898-5862
badler@central.cis.upenn.edu

1 Summary

This project adapted the *Jack* human modeling system to be ready to interface to emerging Virtual Reality systems.

2 Results

The principal results of this project were internal changes to the *Jack* human modeling system to enable support of external Virtual Reality devices. In particular, we improved the human hand model in order to ready it for input from glove devices such as the VPL *DataGlove*®, and obtained an Ascension Technologies *Flock of Birds*®for multiple 6 degree-of-freedom data inputs.

In the latter case, the *Flock of Birds* device served as the input for whole body motion from just 4 sensors: one on each hand, one on the head, and one on the base of the spine. The hand sensors provided end effector position and orientation; the head sensor provided a reasonable approximation to eye gaze, the spine sensor gave the approximate position and rotation of the center of mass in the ground plane, and the combination of head and spine sensors gave the torso attitute [1]. Thes inputs were configured to drive constraints, which were then solved in real-time by our inverse kinematics algorithm and stepping behaviors embedded in our *Jack*®software citeBPW93. The results were very encouraging, and led to the adoptation of the scheme by other groups (such as Sandia National Labs) for avatar motion input.

References

[1] Badler, N., Hollick, M. and Granieri, J. Real-Time Control of a Virtual Human using Minimal Sensors. *Presence* 2(1), 1993, pp. 82-86.

[2] Badler, N., Phillips, C. and Webber, B. *Simulating Humans: Computer Graphics, Animation and Control.* New York: Oxford University Press, 1993.

ARO Workshop on Integrating Heterogeneous Information Systems
Final Report – Subcontract to DAAL03-89-C-0031

Peter Buneman
Department of Computer & Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389

ARO - AI CENTER WORKSHOP ANNOUNCEMENT
ARO Workshop on Integrating Heterogeneous Information Systems

Location: Naval Postgraduate School, Monterey, CA

Date: March 8 and 9, 1993

Organizers: Peter Buneman, Susan Davidson (University of Pennsylvania)
Sham Navathe (Georgia Institute of Technology)
Magdi Kamel (Naval Postgraduate School)

Workshop Focus: Given the increasing need to integrate previously isolated or independently developed information systems, the subjects of "Heterogeneous Databases", "Enterprise Integration" and "Concurrent Engineering" have recently become the focus of considerable research, and some interesting prototype integration tools and methods have been developed. Although there have been abundant workshops in each of these subjects, areas of commonality between these subjects has not been explored. The purpose of this workshop is to bring together academic and industrial researchers in heterogeneous information systems, enterprise integration and concurrent engineering to explain the fundamental problems and explore areas of overlap and application. We will also explore an application in these problems, that of Biological Databases (the Genome project), and the BrainMap database.

The workshop is sponsored by the Army Research Office, and is open free of charge to anyone affiliated with the Army.

Schedule

Monday, March 8:

9:00-9:15	Introductions and General Information (Susan Davidson, Magdi Kamel)
9:15-12:30	Sample Problem Areas of Integration (Peter Buneman, moderator): Chris Overton, *In Biological Databases, the Problem is Merging Conceptual Schemas* Sham Navathe, *An overview of the FADB Project* Dan Fishman, *A Framework for Enterprise Integration*
12:30-2:00	Lunch (on your own)
2:00-4:00	Engineering Integration and Systems (Sham Navathe, moderator): Sudha Ram, *Information Sharing among Heterogeneous Databases in a Manufacturing Environment* Martin Hardwick, *Implementing Concurrent Engineering using PDES/STEP and Information Modeling* David Hsiao, *The Multi-backend, Multi-model, Multi-lingual System*
4:00-4:30	Coffee Break
4:30-6:00	Biological Database Integration (Victor Markowitz, moderator): Peter Fox, *BrainMap: A database for visualizing human functional neuroanatomy* Peter Karp, *Towards the Biomatrix*
6:00	Reception

Tuesday, March 9:

9:00-12:30	Techniques for Integrated Information Retrieval (Chris Overton, moderator): Peter Buneman, *Database query languages are more powerful than you thought* Hector Garcia, *TSIMMIS: An Integrated Information Management System* Susan Davidson, *Federating Approximations*
12:30-2:00	Lunch (on your own)
2:00-4:00	Tools and Interfaces for Integrating Genomic Data (Hector Garcia, moderator): Magdi and Nabil Kamel, *An Integrated Environment for Molecular Biology Databases* Victor Markowitz, *The Integrated Information Management System for Human Genome Data* Manfred Zorn, *Automatic Generation of User Interfaces and Implications for Database Integration*
4:00-4:30	Coffee Break
4:30-5:30	Discussion and Wrap-up (Peter Buneman and Susan Davidson)

Titles and Abstracts of Talks

Database query languages are more powerful than you thought

Peter Buneman (University of Pennsylvania)

Database query languages have a particularly simple structure. The reason for this is partly to make them simple to use, but more importantly to make it possible to optimize them by program transformation. The traditional example of this is the relational calculus (and SQL-like languages) which can be optimized by transformations of the relational algebra. Recently, a much more general algebra has been found to describe languages that resemble the relational calculus. It is more general (a) because it can deal with nested structures, such as relations that contain relations, and (b) because it can equally well describe queries on other "collection" types such as multisets and lists.

I shall describe the basic structure of this language and show how it may be used in querying "non-traditional" databases such as ASN.1 data and gene data described by grammars. Time permitting, I shall also suggest that the language may also be appropriate for numerical computations. Because of this generality, I believe the next generation of query languages may provide very powerful interfaces for heterogeneous database *and* to "non-database" structures.

A Framework for Enterprise Integration

Dan Fishman (Hewlett-Packard Labs)

Several different models are being pursued as ways to approach enterprise integration. These models are based on different views of the core activities of the enterprise. In my presentation, I will discuss three such views:

- the OLTP-centric view which aims to ensure that operational systems interoperate effectively and efficiently in the execution of business transactions;
- the distributed object computing or CORBA-centric view, which aims to provide a distributed object messaging infrastructure whereby objects may invoke the services of other objects in a location-transparent manner;
- and the information-centric view which aims to provide an integrated, single logical view of enterprise information and a medium through which this information is accessed and used.

Each of these views plays a useful role in integrating the enterprise, but each addresses only part of the problem. Thus, I will also present a framework that combines these and other views to address the enterprise integration problem more broadly.

BrainMap: A database for visualizing human functional neuroanatomy

Peter Fox, MD (Research Imaging Center University of Texas Health Science Center at San Antonio)

The Human BrainMap Database (BrainMap) visualizes the functional organization of the human brain. Knowledge of the functional architecture of the human brain is expanding rapidly as the use of the non-invasive, brain-imaging methods spreads. Centrally maintained as a cumulative record of brain-mapping research, BrainMap allows scientists to contribute observations and to keep abreast of progress in laboratories world-wide. Visualization of the brain's functional architecture is compute and data intensive, due both to the graphical detail needed for anatomical visualization and to the information about functionality which must be retrieved, overlayed and visualized. Additionally, transmission of bulky image data over internet is extremely slow. BrainMap has resolved these difficulties by two means. First, all graphics reside locally as a digitized atlas of brain photographs

and outlines. Only ascii data are retrieved from the central server or transmitted via internet. Second, to allow the use of a common set of graphics, all data are coded for brain location using a Cartesian coordinate system normalized for variations in brain size and shape. Shape-normalized anatomical coordinates have become an international standard for reporting human-brain functional organization. Optimizing graphics and database management independently, this strategy achieves: 1) rapid, high resolution visualization, 2) rapid retrieval and transmission of anatomically specific functional information, 3) multi-site data entry, and 4) the managerial simplicity of a centralized database. We offer BrainMap as a model for databases requiring visualization of complex structures.

TSIMMIS: An Integrated Information Management System
Hector Garcia-Molina (Stanford University)
TSIMMIS is a joint Stanford/IBM Almaden project whose goal to develop a next generation information management system that provides integrated access to a wide variety of information and knowledge sources. The system will manage or provide access to unstructured text data, such as newswire stories and files, to conventional structured databases, to raw data from satellite images and simulation experiments, and to knowledge bases containing rules for information processing. Through this system a decision maker will be able to find information of interest, fuse information from different sources, and process it (e.g., summarize it, visualize it, discover trends).
In this talk I will focus on one particular component of the project: constraint and rule management across heterogeneous sites. Even if information sites are autonomous and heterogeneous, then may need to guarantee certain constraints (example: the jet engine designed in one system by one company must be "attachable" to the wing designed and build by a different company). I will briefly discuss how such constraints can be specified, how they can be enforced, and what guarantees can be made about them. (Joint work with Sudarshan S. Chawathe and Jennifer Widom.)

Implementing Concurrent Engineering using PDES/STEP, and Information Modeling
Martin Hardwick (RPI & STEP Tools, Inc.)
For the last 10 years a community of researchers and practitioners have been developing system independent information models for engineering data in an effort known as PDES in the US, and STEP in the rest of the world. The result of these efforts will be published in 1993 as the ISO 10303 STEP standard. In this talk we will describe how the PDES/STEP information models have been used to implement concurrent engineering at various pilot sites as part of the DARPA Initiative in Concurrent Engineering (DICE). Some of the issues that will be discussed include translating information models into data models, and using delta files to coordinate engineering changes.

The Multimodel, Multilingual, and Multibackend Database System
David K. Hsiao (Naval Postgraduate School)
The multimodel, multilingual, and multibackend database system is a solution towards data sharing and resource consolidation of heterogeneous databases and systems. By multimodel, it is meant that the user can create a database in the data model of user's choice. At the present, the relational user can create relational databases, the hierarchical user hierarchical databases, the network user network databases, object-oriented user object-oriented databases all in the same system. By multilingual, it is meant that the user can write transactions in the user's favorite data language for the purpose of accessing and manipulating not only the user's homogeneous databases, but also other users' heterogeneous databases. For example, a relational user can write SQL transactions for the purpose of accessing and manipulating not only relational databases, but also hierarchical and object-oriented databases as if they are relational. By multibackend, it is meant that all these

heterogeneous databases and system software are supported by a single database computer with a scalable number of parallel backends. The multimodel and multilingual features are intended for data sharing and exchange among heterogeneous databases. The multibackend feature is intended for resource consolidation of all the database store, system software, computer hardware, and personnel support.
In the Laboratory for Database Systems Research, an experimental multimodel, multilingual, and multibackend database system can demonstrate the aforementioned features.

An Integrated Environment for Molecular Biology Databases
Magdi Kamel (Naval Postgraduate School), Nabil Kamel (University of Florida)
The requirements of a traditional integration of databases include: hiding the heterogeneity among its member databases, preserving their autonomy, supporting controlled data sharing, and providing a user- friendly interface, with a performance comparable to that of a homogeneous distributed database system. In this talk, we describe a variation of the loosely-coupled federated database approach that satisfies these requirements and is specifically oriented towards the special characteristics and needs of molecular biology databases. The approach differs from the traditional federated database approach in that it also integrates software tools together with the databases. The talk will also include a description of XBio, a system which constructs such an integration.

Towards the Biomatrix
Peter D. Karp (AI Center, SRI International)
Consider the Biomatrix to be an interconnected collection of biological databases and knowledge bases that are to be used for a variety of queries and complex computations. This talk will discuss some of the design objectives and constraints for the Biomatrix, and will discuss some of the anticipated problems in connecting existing biological databases based on the properties of those databases. I propose that it is difficult to envision creating a useful Biomatrix that lacks some sort of global schema, and posit that AI knowledge representation techniques provide a solid foundation on which to build a global schema that unifies the diverse semantics of many existing databases.

The Integrated Information Management System for Human Genome Data
Victor Markowitz (Lawrence-Berkeley Labs)
There are currently numerous experimental databases as well as several major databases containing data pertaining to the Human Genome project. These databases provide a variety of information including nucleic acid sequences (e.g. from GenBank, EMBL), peptide sequences (e.g. PIR, SwissProt) as well as mapping data (e.g. GDB, CEPH). These databases contain related or even overlapping data, and are characterized by various degrees of heterogeneity: they reside on different hardware under different operating systems, are based on different DBMSs, and reflect different views of data. Mechanisms assisting scientists in resolving various heterogeneity problems are essential for allowing them to retrieve and correlate data from multiple genomic databases. No such mechanisms are currently available.
We will discuss some aspects of developing the Integrated Information Management System for Human Genome Data (IGD). IGD will provide scientists with a unified view of multiple genomic databases by collecting data of interest into a comprehensive back-end database. IGD will provide tools for retrieving, displaying, analyzing and editing data on local front-ends, and interfaces to existing analytical tools. Furthermore, IGD will support the integration of data with biological knowledge.
The IGD project is a collaborative effort involving groups at the German Cancer Research Center in Heidelberg, at MRC Laboratory of Molecular Biology in Cambridge, England, at Imperial Cancer

Research Fund in London, at CNRS-CRBM in Montpellier, France, and at Lawrence Berkeley Laboratory.

An overview of the FADB (Federated Autonomous Databases) Project with Honeywell, U. of Minnesota and Georgia Tech.

Sham Navathe (Georgia Tech)

This project addresses the understanding and solving of the problem of providing integrated access to a variety of autonomous databases. While we assume that the databases are heterogeneous and distributed, in our current effort we have used the relational model with SQL2/3 as the basis for our work. The presentation will describe the proposed overall architecture that involves a set of "component databases" to be shared, the resulting views called "federated databases", gateways, and client applications. The topics researched so far consisted of schema integration, federated query processing, and consistency management. The presentation will highlight the work on "instance-level integration" and entity identification, multi-phased query optimization, and some results from concurrency control and recovery. The work has been supported by Rome Labs of the U.S. Air Force.

In Biological Databases, the Problem is Merging Conceptual Schemas

Chris Overton (Dept of Genetics, University of Pennsylvania School of Medicine)

Computational biology depends on ready access to information contained in multiple databases such as the nucleic acid sequence (DDBJ, EMBL, GenBank), protein sequence (PIR, SwissProt), biological macromolecular structure (PDB), genome (GDB, Mouse Encyclopedia, ACEDB) databases along with bibliographic information (Medline) and numerous even more highly specialized databases. Querying and extracting information from across several of these databases is a daunting task, requiring an intimate knowledge of the conceptual structure of each database and how concepts relate across databases. In this talk, I will describe the structure of the databases in common use in our laboratory – SORTEZ, a relational version of the ASN.1 schema for GenBank, GDB, a relational database of human genome information, TFDB, a Prolog version of a transcription factor database, and Chr22DB, a local relational database for a Human Genome Center laboratory notebook database – and our efforts to integrate them through a Prolog interface. In our experience, the major impediment to multiple database access is schema integration, and thus from a purely pragmatic point of view, our most compelling need is for tools that aid in the integration process.

Information Sharing among Heterogeneous Databases in a Manufacturing Environment

Sudha Ram (Department of Management Information Systems, University of Arizona)

This research deals with the integration of heterogeneous databases and file systems in a manufacturing environment. We will report on some of our experiences in providing access to the information base for the Automated Manufacturing Research Facility (AMRF) at the National Institute of Standards and Technology. The Automated Manufacturing Research Facility (AMRF) at the National Institute of Standards and Technology (NIST), funded by NIST and the Navy Manufacturing Technology Program, is a laboratory providing a basic array of discrete metal parts manufacturing equipment and control systems - a testbed for identification and exercise of potential standard interfaces between existing and future components of small-batch manufacturing systems. In addition, the AMRF provides a laboratory for the development of factory-floor metrology in an automated environment and a testbed for research on the next generation of "knowledge-based" manufacturing systems. To provide a realistic testbed for interface standards, the AMRF is intentionally composed of manufacturing, computing equipment and data management software from

many vendors, comprising over 20 individual control systems, all of which are interconnected over a common network, along with the manufacturing engineering systems, including computer-aided-design (CAD) systems and process planning systems. This presentation will describe our efforts at achieving interoperability among the heterogeneous systems at the AMRF.

Automatic Generation of User Interfaces and Implications for Database Integration
Manfred Zorn (Human Genome Computing Group, Lawrence Berkeley Laboratory)
Databases for genomic data are subject to continuing evolution to cope with scientific advances. Modifications in the database definition invariably trigger changes in the user interface. Thus a significant effort is spent in constantly adapting the database to new requirements and the user interface to the modified database definition. To break this vicious cycle we have developed a user interface that is automatically generated from the database definition, i.e., the metadata.
A generic user interface guides the user through a standard flow of actions, from object selection and query formation to viewing the details of an object and following connections to linked objects. A plain text configuration file read upon program start-up provides information for a specific database, e.g., names of objects, attributes, labels, thus customizing the generic interface for a particular application.
We will describe how this approach can be extended to allow dynamic adaptations of a user interface to a variety of databases and discuss the problems with current genomic databases.

Final Report
ARO AI Scholars Summer Institute
(held May 20 – June 12, 1992)

Gregory Provan, PI

June 30, 1992

The Second Annual ARO AI Scholars Summer Institute was a three and one half week summer research experience for a group of 12 undergraduate, under-represented minority students in the research areas falling under the scope of the ARO-sponsored AI Center of Excellence — including natural language processing, computer graphics, and robotics. It was intended that the students would be chosen from historically minority institutions, and would be either junior or seniors in engineering studies. We proposed an intellectual program of daily intensive laboratory research work and training in programming and computation in artificial intelligence, with extensive interaction with leading faculty members and graduate students, who would act as mentors for the undergraduates, and appropriate social activities planned.

1 The Program

With funding from the Army Research Office, the Comprehensive Regional Center for Minorities (CRCM), Philadelphia Partnership for Education, and the Institute for Research in Cognitive Science (IRCS), the second ARO AI Scholars Summer Institute was held from May 20 – June 12, 1992. Advertisements of the Summer Institute were posted at 20 colleges and universities, both historically black colleges and other universities. Twelve minority students were selected upon the basis of an application consisting of a statement of goals for a future career, an official transcript, along with school and major subject information, and a letter of recommendation. The students were given weekly stipends, housed in a University of Pennsylvania dormitory, and were provided with a University meal plan. Two African-American male graduate students from the School of Engineering and Applied Sciences (SEAS) lived in the dormitory with the students for the duration of the Institute.

An intensive daily schedule provided the students with a broad range of information in several areas. One week focused on a mini-course on Probability and Statistics in Computer Science and Engineering. The morning sessions were taught by Dr. Max Mintz, and the afternoon *Mathematica* lab sessions were instructed by a CIS graduate student. Other areas of study included Natural Language Processing, Medical Decision Making, Systems Programming, Computer Networks and

Databases/Algorithms.

The students attended lab sessions between 2 and 4 hours per day. These lab sessions related to the morning's topic and were partially coordinated by CIS graduate students. For most of the students this was their first introduction to C Programming. In addition to the topics presented in lecture and lab sessions, the students worked on a programming project of interest for the three weeks of the Institute. Some of the projects included: [describe projects here, possibly]

Various social activities were scheduled — an afternoon at the Franklin Institute, a 2-hour tour of the University Museum, and outings to the New Jersey shore, Great Adventure, as well as activities within Philadelphia.

2 Results

This was the second attempt on the part of the AI Center to reach out to the under-represented minorities community in a deliberate, concentrated way. We consider this attempt to have been an enjoyable and challenging experience for the students, as well as a success in encouraging them to pursue further education in computer science and engineering. Please see Appendix B for feedback we have received from the Scholars.

Most of the students felt that "it was an excellent opportunity to share and get every detail of the graduate studies experience." In addition to exposing the students to state of the art facilities that they did not have at their home institution, it showed them "that large schools are not too intimidating", and that they would now consider applying to schools for graduate studies that they had not considered before.

Many of the comments focused on Professor Mintz's mini-course, and the fact that they were glad to have the opportunity to learn *Mathematica*. The lab sessions challenged many of the students, as they were not accustomed to the concentrated pace at their home school. Many of the students now feel that graduate school is a definite goal to plan for in their future, and in fact are planning to further investigate some of the areas of artificial intelligence that were presented at the Institute.

The Summer Institute also received wide media interest. National Public Radio featured a five minute segment on the Institute on their Morning Edition program. The Philadelphia Daily News had an article on it in June, as did The Daily Pennsylvanian and The Compass, both Penn publications. (Please find copies of these articles Appendix A.)

3 Changes Made to the Program

Most of the recommendations made in last year's program final report were instituted in this year's program.

A formal application process was instituted with applications due April 15. Each application consisted of a statement of goals for a future career, an official transcript and a letter of recommendation in addition to the usual information on school, major, etc.

The program was lengthened to just short of four weeks, and a focus provided both for the faculty lectures as a whole and for the programming laboratories. The lectures were focused around the applications of probability and statistics in Engineering and Computer Science. The main focus of the lecture series was a "mini-course" in probability and statistics given by Prof. Mintz. Many of the other faculty lectures were on applications of probability and statistics in Engineering and Computer Science.

4 Recommendations for Future Programs

Given the expansion of the program to the current level of 12 students for 4 weeks, some critical factors can be inferred from experience. Some rethinking of commitment and objectives is necessary if the program is to be continued at the current level. Some recommendations for expanding the program are also summarized.

4.1 Maintaining Current Level

First, the current level of ARO support for the program is inadequate. Programs of this nature are *inherently* labor-intensive, and significantly reducing the contact-hours of such a program would reduce its quality. To maintain the quality, and hence guarantee the success, of such a program, adequate funding needs to be made available.

A detailed examination of the budget for the program shows that, for the actual number of contact-hours, the cost is relatively low. This is because most of the contact is made by a small number of faculty (who were merely paid honorariums) and graduate students (who were paid at graduate student, i.e. sub-professional, rates). If such a program were staffed and paid at professional rates, then the cost would be significantly higher.

Second, the departmental support provided for the program needs to be re-evaluated. This past year, there were significant under-estimates of the amount of time taken to organize and run such a program. It is an exceedingly time-consuming process to organize a program in which 12 college students are to receive teaching for 8 hours per day for 4 weeks. Since this was the first time that the program was run on this scale, the significant amount of effort in planning the program may not need to be duplicated. Nonetheless, two recommendations need to be considered:

- At the very minimum, the program needs to be run by a faculty director plus an experienced assistant working full-time on the program during part of the planning phases and the program

itself. In addition, an administrator (of the calibre of Shirley Aderman, Karen Carter or Trisha Yannuzzi, senior administrative staff) is needed to arrange for additional staff as requested, both before and during the program. Release from other duties for a secretary/administrator must also be arranged ahead of time and paid for explicitly. The Institute also received strong support from the Business Office, and their effort should also be taken into account, possibly with special arrangements to ensure their support when necessary.

- The faculty compensation needs to be fixed *a priori* on an hourly basis. For the faculty director, the significant amount of time required to run such a program should be planned for *a priori*, and appropriate compensation provided. The time commitment involved in running the program this year was 9 weeks full-time (a conservative estimate). This includes 6 weeks of preparation and the 4 weeks of the program. Clearly, the preparation time will decrease as the program becomes more established, but a minimum of 6-8 weeks effort (possibly half-time) is needed to organize and run the program.

  For the faculty lecturers, adequate compensation levels also need to be established. The current honorarium system is not viable in the long term. For example, the extraordinary effort put out for the program by certain faculty (like Prof. Mintz in teaching a one-week mini-course) will not be repeated unless appropriate compensation is planned. The mini-course was highly successful, and in fact the students requested that more courses of this type be given in future. If we are to do this, we must compensate the relevant faculty appropriately.

5 Expansion of the Program

Following the success of this program, it is not unreasonable to repeat the program with minor modifications, i.e. retain the 4-week, 12-student structure. The present structure is clearly viable, and certainly accomplishes the desired goals. However, possible approaches to expanding the program include:

- Expanding the class size. Increasing the class size to 16 is a relatively simple change. However, increasing the class beyond that may degrade the quality of the program in terms of personal contact-time per student.
- Increasing the length of the program. Increasing the length of the program will allow more in-depth mini-courses and programming projects. College credits could also be given for attending the program. The appropriate length of time would be 6 or 8 weeks.

  However, such an increase in length would also involve more planning and administration effort on the part of the course director, faculty and administrative staff.

6 The Scholars

The Scholars (7 men and 5 women) were computer science or engineering majors between their junior and senior years of undergraduate study. (One student was a sophomore.) Ten of the

students were African American and 2 students were Puerto Rican. The students came from a variety of backgrounds. However, all displayed a desire to get learn more about graduate studies, as can be seen by their activities on their application forms, and their answers on the evaluation forms.

The experience and knowledge of the students varied widely, as did their interests in specific areas of Artificial Intelligence.

1. Juan M. Vazquez Berrios
   Junior
   School: University of Puerto Rico, Mayaguez, Puerto Rico
   Home: Mayaguez, Puerto Rico
2. Vidette Cornelius
   Junior
   School: Lincoln University, Lincoln University, Pennsylvania
   Home: Philadelphia, Pennsylvania
3. Willonda Fisher
   Senior
   School: Cheyney University of Pennsylvania, Cheyney, Pennsylvania
   Home: Philadelphia, Pennsylvania
4. Krista M. Holmes
   Junior
   School: Florida A&M University, Tallahassee, Florida
   Home: Chicago, Illinois
5. Sean K. Johnson
   Junior
   School: Morehouse College, Atlanta, Georgia
   Home: Bloomfield Hills, Michigan
6. D'ondria L. Kennard
   Senior
   School: Prairie View A&M University, Prairie View, Texas
   Home: Fort Worth, Texas
7. Carlos M. Santiago
   Junior
   School: University of Puerto Rico, Mayaguez, Puerto Rico
   Home: Mayaguez, Puerto Rico
8. Walter E. Smith IV
   Senior
   School: Prairie View A&M University, Prairie View, Texas
   Home: Dallas, Texas
9. Kendal E. Stephen
   Junior
   School: University of Maryland, Baltimore, Maryland
   Home: Largo, Maryland
10. Lisa M. Stephen
    Sophomore

School: Oakland University, Rochester, Michigan
Home: Detroit, Michigan

11. Gregory J. Toatley, Jr.
    Senior
    School: Lincoln University, Lincoln University, Pennsylvania
    Home: Willingboro, New Jersey
12. Billie R. Whitfield, Jr.
    Senior
    School: University of Maryland, Baltimore, Maryland
    Home: Lanham, Maryland

7 Penn Participants

Below is a list of Penn faculty and students who participated in the Institute. Please refer to the Schedule of Activities shown on pages 8 – 15 for specific descriptions of how each of these individuals participated.

1. Norman Badler, Professor, CIS, and PI.
2. Edward Brown, student mentor, EE graduate.
3. Susan Davidson, Associate Professor, CIS.
4. Ralph Etienne-Cummings, student mentor, EE Graduate Student.
5. Mike Felker, Assistant for Graduate Admissions.
6. Cora Ingrum, Director, SEAS Minorities Engineering Program.
7. Insup Lee, Associate Professor, CIS.
8. Elizabeth Levison, lab instructor, CIS graduate student.
9. Robert Mandelbaum, *Mathematica* lab instructor, CIS graduate student.
10. Mitch Marcus, Professor, CIS.
11. Max Mintz, Associate Professor, CIS.
12. Amarnath Mukherjee, Assistant Professor, CIS.
13. Jeffrey Nimeroff, lab instructor, CIS graduate student.
14. Gregory Provan, Institute Director, Assistant Professor, CIS.
15. Sanguthevar Rajasekaran, Assistant Professor, CIS.
16. Philip Resnik, lab instructor, CIS graduate student.
17. Officer Riley, University of Pennsylvania, Public Safety Officer.
18. Mark Steedman, Associate Professor, CIS, Chair of Graduate Admissions.
19. Many other CIS graduate students participated in the informal group session entitled, "What is Life as a Graduate Student *Really* Like?"

Posters and applications were distributed to the following schools:

1. University of Maryland, Baltimore County
2. Georgia Institute of Technology
3. Morgan State University, Baltimore MD
4. Prairie View A & M University, Texas
5. University of Puerto Rico - Mayaguez
6. Howard University, Washington, DC
7. Lincoln University, PA
8. South Carolina University, Orangeburg, SC
9. North Carolina Agricultural & Technical State University, Greensboro, NC
10. Tennessee State University, Nashville, TN
11. Clark Atlanta University, Atlanta, GA
12. Morehouse College, Atlanta, GA
13. Florida A. & M. University, Tallahassee, FL
14. Cheyney University, Phila.
15. St. Joseph's University, Phila.
16. Hampton University, VA
17. Temple University, Phila.
18. LaSalle University, Phila.
19. Bryn Mawr University, PA
20. Drexel University, Phila.

8 Schedule of Activities

Calendar of Activities
ARO Scholars Program
20 May 1992 – 12 June 1992

Tuesday, May 19:

	Students arrive, Harrison House 3820 Locust Walk
7:00 p.m.	Dinner at Boccie's, 4040 Locust Street (meet in Harrison House lobby, 6:45 p.m.) (Tables reserved in Dr. Provan's name).

Wednesday, May 20:

8:30 – 9:00	Breakfast in Raisler lounge
	Towne Building
9:00 – 9:30	Dr. Farrington (SEAS Dean): welcome speech
9:30 – 10:00	PENN ID Center, Room 323A, 3401
10:00 – 12:00	Nimeroff/Resnik/Levison System Introduction: Review of Unix commands and tools (554 Moore)
12:00 – 1:15	LUNCH
1:30 – 2:30	Officer Riley, Public Safety — Welcome to Penn/Security Guidelines (554 Moore)
3:00 – 5:00	Nimeroff/Levison C Programming Basics (MAC Lab)
5:00 – 6:30	Dinner at 1920 Commons.

Thursday, May 21:

10:00 – 12:00	Badler: Computer graphics and animation (Room 107 Moore)
12:00 – 1:15	LUNCH
1:30 – 3:30	Nimeroff/Levison: C Programming Lab (MAC Lab)
3:30 – 5:00	Lab work in the Graphics Lab (Room 107 Moore)
5:00 – 6:30	Dinner, 1920 Commons.
7:00 – 9:00	Dessert with Michel deGraff, UPenn graduate student

Friday, May 22:

10:00 – 12:00	Badler: Computer graphics and animation (Room 107 Moore)
12:00 – 1:15	LUNCH
1:30 – 3:30	Nimeroff/Levison: C Programming Lab (MAC Lab)
3:30 – 5:00	Lab work in the Graphics Lab (Room 107 Moore)
5:00 – 6:30	Dinner, 1920 Commons.

Saturday, May 23

10:00	Trip to African-American Museum (meet in Harrison House lobby, 9:45)

Week 2 (May 25 – May 29)

MINI-COURSE ON PROBABILITY AND STATISTICS IN COMPUTER SCIENCE AND ENGINEERING

Monday, May 25: HOLIDAY: MEMORIAL DAY

Trip to NJ shore or Great Adventure Amusement Park (TBA)

Tuesday, May 26:

9:30 – 10:30	Mintz: Probability and Statistics for Computer Science and Engineering (Alumni Hall, Towne Building)
10:30 – 12:00	Mintz: Probability and Statistics (Alumni Hall, Towne Building)
12:00 – 1:15	LUNCH
1:30 – 3:30	*MATHEMATICA* LAB (Lab work on Probability and Statistics) Mandelbaum (MAC Lab)
3:30 – 5:00	Nimeroff/Resnik: C Programming Lab
7:30 –	Mixer with Engineering Graduate Students (Rooftop Lounge, Harrison House).

Wednesday, May 27:

9:30 – 10:30	Mintz: Probability and Statistics for Computer Science and Engineering (Alumni Hall, Towne Building)
10:30 – 12:00	Mintz: Probability and Statistics (Alumni Hall, Towne Building)
12:00 – 1:15	LUNCH
1:30 – 3:30	*MATHEMATICA* LAB (Lab work on Probability and Statistics) Mandelbaum (MAC Lab)
3:30 – 5:30	Resnik: C Programming Lab

Thursday, May 28:

9:30 – 10:30	Mintz: Probability and Statistics for Computer Science and Engineering (Alumni Hall, Towne Building)
10:30 – 12:00	Mintz: Probability and Statistics (Alumni Hall, Towne Building)
12:00 – 1:15	LUNCH with Cora Ingrum Director, Minorities Programs - SEAS
1:30 – 3:30	*MATHEMATICA* LAB (Lab work on Probability and Statistics) Mandelbaum (MAC Lab)
3:30 – 5:00	Nimeroff/Levison: C Programming Lab

Friday, May 29:

9:30 – 10:30	Mintz: Probability and Statistics for Computer Science and Engineering (Alumni Hall, Towne Building)
10:30 – 12:00	Mintz: Probability and Statistics (Alumni Hall, Towne Building)
12:00 – 1:15	LUNCH
1:30 – 3:30	*MATHEMATICA* LAB (Lab work on Probability and Statistics) Mandelbaum (MAC Lab)
3:30 – 5:00	Nimeroff/Levison: C Programming Lab

Saturday, May 30:

Self-guided tour of Center City Philadelphia.

Sunday, May 31:

6:00 a.m. – 10:00 p.m.	Trip to Wildwood and Cape May Meet in Harrison House lobby at 6:00 a.m.

Week 3 (June 1 – June 5)

APPLICATIONS OF PROBABILITY AND STATISTICS IN COMPUTER SCIENCE AND ENGINEERING

Monday, June 1 (Natural Language Processing):

9:30 – 10:30	Marcus: Statistics and Natural Language Processing (554 Moore)
10:45 – 12:00	Resnik: Introduction to Programming project (Room 302C, 3401 Walnut St.)
12:00 – 1:15	LUNCH
1:30 – 2:45	Nimeroff: Parallel processing and introduction to the Connection Machine (CM2) (Suite 3C, 3401 Walnut St.)
3:00 – 5:00	Nimeroff/Resnik: Project Lab (MAC Lab)
5:30 – 6:30	Dinner in 1920 Commons

Tuesday, June 2 (Natural Language Processing):

9:30 – 10:30	Marcus: Statistics and Natural Language Processing (554 Moore)
10:45 – 12:00	Mark Steedman: Natural Language Understanding (554 Moore)
12:00 – 1:15	LUNCH
1:30 – 2:45	Nimeroff/Levison: C Programming Techniques: Makefiles (MAC Lab)
3:00 – 5:00	Nimeroff/Levison: Project Lab (MAC Lab)
5:30 – 6:30	Dinner in 1920 Commons

Wednesday, June 3:

9:30 – 10:30	Etienne-Cummings: Computer hardware for neural computing (Room 554 Moore)
10:45 – 12:00	Demonstrations: EE Lab, GRASP Lab
12:00 – 1:15	LUNCH
1:30 – 2:45	Resnik: Project Lab (MAC Lab)
3:00 – 5:00	Resnik: Project Lab (MAC Lab)
5:30 – 6:30	Dinner in 1920 Commons

Thursday, June 4 (Medical Decision Making):

9:30 – 10:30	Provan: Medical Decision Making (Room 318C, 3401 Walnut St.)
10:45 – 12:00	Nimeroff: Programming tools for medical decision making (Room 318C, 3401 Walnut St.)
12:00 – 1:15	LUNCH
1:00 – 1:15	Meet at Moore School Parking Lot
1:30 – 5:00	Visit to the Franklin Institute
5:30 – 6:30	Dinner in 1920 Commons

Friday, June 5:

9:00 – 11:00	Nimeroff/Resnik: C Programming Lab (MAC Lab)
11:00 – 12:00	Mukherjee: Probability and networking (Room 318C, 3401 Walnut St.)
12:00 – 1:00	LUNCH
1:30 – 2:45	Nimeroff/Resnik: Project Lab (Suite 3C, 3401 Walnut St.)
3:00 – 5:00	Nimeroff/Resnik: Project Lab (MAC Lab)
5:30 – 6:30	Dinner in 1920 Commons

Week 4 (June 8 – June 12)

PROJECT PROGRAMMING AND DATABASES/ALGORITHMS

Monday, June 8:

Ms. Veniece Keene, CRCM Program Officer is visiting today.

9:30 – 10:00	Mark Steedman, Professor Chair of Graduate Admissions Mike Felker, Asst. for Graduate Admissions (554 Moore)
10:00 – 12:00	Davidson: Databases (554 Moore)
12:00 – 1:00	LUNCH
1:00 – 3:00	Lab work on Databases (MAC Lab)
3:00 – 5:00	Nimeroff/Resnik: Project (MAC Lab)
5:30 – 6:30	Dinner at 1920 Commons

Tuesday, June 9:

10:00 – 11:30	Lee: Introduction to concurrency and process algebra (554 Moore)
11:30 – 1:00	LUNCH
1:15 – 2:45	Nimeroff/Resnik: Project (MAC Lab)
3:00 – 5:00	Resnik: Programming Project (Room 302C, 3401 Walnut St.)
5:00 – 6:30	Dinner at 1920 Commons

Wednesday, June 10:

10:00 – 11:30	Lee: Introduction to concurrency and process algebra (554 Moore)
11:30 – 1:15	LUNCH
1:30 – 3:00	Tour of University Museum (Meet in Education Office near Kress Entrance of the Museum).
3:30 – 5:00	Project Wrap-up (3401 Walnut, Room 318C)
5:30 – 6:30	Dinner at 1920 Commons

Thursday, June 11:

9:30 – 10:30	Rajasekaran: Algorithms (3401 Walnut, Room 318C)
10:30 – 12:00	Visit to Graphics Lab (Room 106 Moore)
12:00 – 1:15	LUNCH
1:30 – 3:00	Rajasekaran: Algorithms (3401 Walnut, Room 318C)
3:00 – 4:30	Q/A Session with CIS Grad. Students What is life as a graduate student *really* like? (Room 554 Moore)
5:00 – 6:30	Dinner at 1920 Commons

Friday, June 12:

10:00 – 12:00	Summer Institute Wrap-up (3401 Walnut, Room 318C)
12:00 – 2:00	Lunch and Presentation of Certificates
2:00 – 5:00	FREE

Saturday, June 13:

9:00 – 11:00	Check out of Harrison House

A Newspaper Articles

The Daily Pennsylvanian,
University of Pennsylvania,
6/14/92

Future engineers visit U.

By GEMINI WAHHAJ
Summer Pennsylvanian Staff Writer

Vidette Cornelius used to have some doubts about getting a graduate degree in engineering.

But that was two weeks ago.

Cornelius, a Lincoln University senior, said her outlook changed after spending the past two weeks participating in a University program that encourages minority students to pursue advanced science degrees.

"I had no idea what to expect in grad school," she said this week. "But now I think, 'Yeah, I can do it.' "

Please see INSTITUTE, page 4

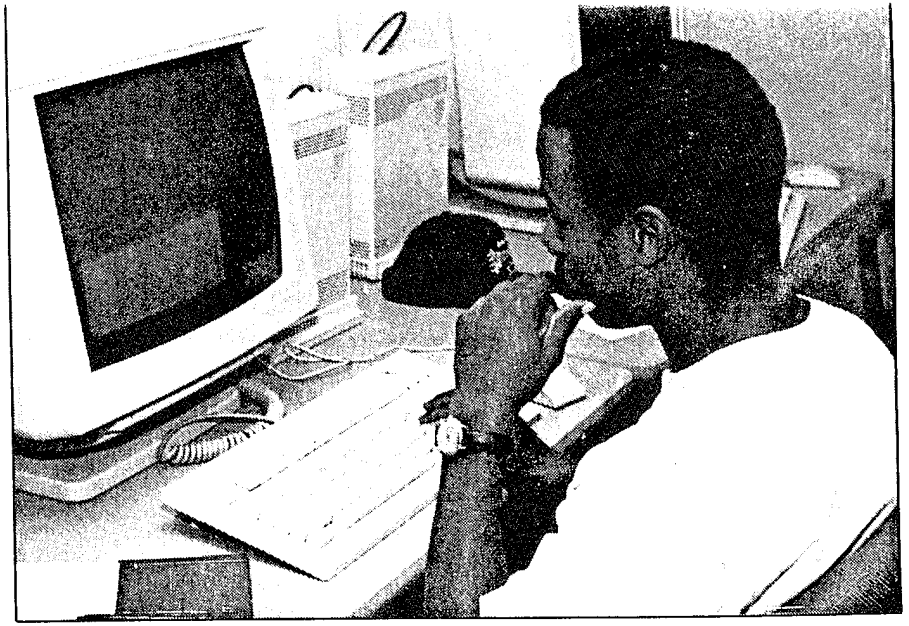

LEANNE CHUKOSKIE/SP Photographer

A student works in a University computer lab as part of the Summer Minority Institute's program to attract more minorities to the sciences.

Engineer hopefuls get a boost during U.'s program for minorities

INSTITUTE, from page 1

Cornelius attributes her new attitude to the program, known as the Summer Minority Institute, which attracted twelve undergraduates from predominantly minority colleges stretching from Michigan to Puerto Rico.

The program's participants have toiled away in the Moore Building's computer labs — and considered it a rare privilege. Organizers said they hope the junior and senior engineering majors will now consider graduate school a more realistic option for themselves and other minority students.

The Institute features workshops focusing on computer applications such as programming, graphics, artificial intelligence and the use of technology in everything from medical decision making to the study of the English language.

Program coordinator Greg Provan, an assistant professor of computer and information sciences, said students seem impressed by the demonstrations.

"Now that they know more about computers, they don't have to be intimidated by these things," Provan said.

Besides taking classes in computer applications, participants have toured workstations and research labs, met with engineering faculty and graduate students, and travelled to the Franklin Institute and the Jersey Shore.

Students said the program has provided a wealth of information about graduate education in the engineering field, and has also reassured them.

D'ondria Kennard, a senior at Prairie View A&M University, said she had been unsure about graduate school until the program exposed her to new technology unavailable in her school's labs.

"We were academically prepared but technologically behind," she said, adding that the economy's sluggish recovery has made graduate school an even more attractive option.

Morehouse College junior Sean Johnson said that talking with graduate students allowed him "to see what life was really like in graduate school."

Willonda Fisher, a senior at Cheney University, said it is important that people appreciate the extra support minority students need to flourish in an area where there are not many role models.

Still, some students remain apprehensive about spending four years in graduate school when they could be earning money in the workplace.

"It's a big decision for a minority student," Johnson said.

The program's participants have toiled away i the Moore Building's computer labs — and co sidered it a rare privilege. Organizers said the hope the junior and senior engineering majors will now consider graduate school a more realistic option for themselves and other mino ity students.

Nationwide, minority students make up fewer than 10 percent of the graduate student population and five percent of science and engineering faculty. There is currently only one black faculty member on the Engineering School's faculty and about three blacks receive engineering PhD's each year.

Students were accepted to the program based on their academic standing, faculty recommendatio and a short essay, Provan said. I added that organizers picked s dents "who'd have the potential go to graduate school and would interested in doing that."

Funding for the program was p vided by the Army Research Offi the Philadelphia Partnership Education, and the University's stitute for Research on Cognit Sciences.

The Compass,
University of Pennsylvania,
Vol. 3, No. 18, 6/18/92

Program for minority students promotes science and engineering

Undergraduate science and engineering students can tell you how sand is made into glass. Graduate students in those fields work at more sophisticated levels and learn how to manipulate matter into high-technology products such as supercomputers and artificial hearts.

It is the knowledge and skill derived from graduate degrees that provide the scientific expertise necessary for technological breakthroughs. According to several national studies, however, not enough Americans—especially from minoritiy groups—are pursuing graduate study and careers in academic research.

Penn's School of Engineering and Applied Science, however, has taken a step toward addressing the situation. During the past several weeks, 12 minority students who have shown a strong aptitude in science and engineering got a first-hand look at graduate education and research through the programs SEAS offers.

Starting on May 19, the students, mostly from historically black colleges, lived and worked at Penn, studying computer graphics, computer languages, mechanical engineering, and other high-technology fields. The program, which ended on June 12, was primarily funded by a grant from the Army Research Office, with additional funding from the Philadelphia Partnership for Education (Comprehensive Regional Center for Minorities) and the University of Pennsylvania's Institute for Research in Cognitive Science.

Not all of the program time was spent in classrooms: there were trips to the Jersey shore as well as other cultural and recreational activities.

The students applied for the program, sending in transcripts, written recommendations, and a short essay on why they wanted to participate. According to Dr. Gregory Provan, an assistant professor of computer and information science at Penn who coordinated the program, the students who were selected met guidelines from the Army Research Office regarding their intent to pursue scientific careers.

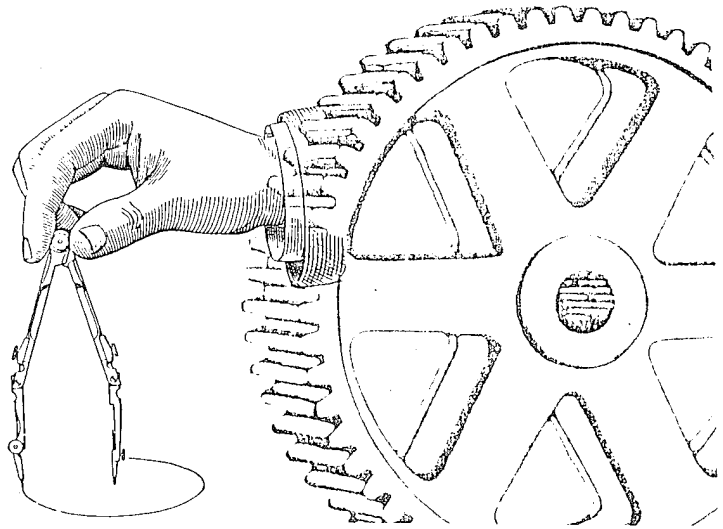

Looking for aptitude—and enthusiasm

"I wanted to see if they were interested in careers in science and not just making money with industry," says Provan. "And I wanted to see *enthusiasm*—that they really wanted to come here rather than just do something for the summer."

Provan, who is African-American, also hopes that working with him and other minority graduate students at Penn showed the students that they, too, can enter and finish graduate school.

In 1991, less than four percent of all Ph.D. degrees in science and engineering in the United States went to members of minorities, and less than five percent of faculty members specializing in science and engineering at colleges and universities were minorities.

The students who took part in the summer program came from such colleges as the University of Puerto Rico; Lincoln University and Cheyney University, both in Pennsylvania; Florida A&M University; Morehouse College in Atlanta; Prairie View A&M University in Texas; the University of Maryland; and Oakland University in Michigan.

—Jon Caroulis

people, places, things. . .

The Philadelphia Tribune,
June 2, 1992

Black music celebration

ASALA, in conjunction with the Black Administrators, faculty and staff of Drexel University present: "A Celebration of Black Music" featuring: Overbrook High School Jazz Band W.C. Bryant School Choir and special guest performers, June 7, 3:30-5:30 P.M., Drexel University, Main Auditorium, 32nd and Chestnut streets. Proceeds benefit Overbrook High School Music Department, William Ross Scholarship Fund, and the choirs participating.

The scoop on summer ed

Twelve students, mostly juniors from historically Black colleges majoring in science and technology, are going to spend five weeks living and learning at Penn. They'll be getting a taste of graduate school studies and research. The program begins May 22 and runs to June 12. Its goal is to encourage non-whites to pursue graduate school and possible careers in academia.

There will be courses in computer science featuring graphics and language, and mechanical and electrical engineering. One key is that they'll be studying with senior faculty (for instance, Norman Badler, who teaches computer graphics, is one of the top people in the world in his field).

In addition, two current Penn students, a senior and a grad student, will act as mentors, helping teach and guide the visiting students. There will also be recreational activities, such as an outing to a Phillies game and a trip to the Jersey shore.

Picasso summers

"Picasso and Things: The Still Life of Picasso" opens Sunday, June 7, at the Philadelphia Museum of Art for an exclusive East Coast showing before heading on to Paris, and the Four Seasons Hotel has a special offer that includes VIP passes to this landmark exhibition.

For the duration of the exhibition (June 7-Aug. 23), the hotel will offer "Picasso '92." The package includes a Four Seasons guest room and two "Anytime" passes to "Picasso and Things." These tickets allow the holders to enter the exhibition whenever they wish during museum hours rather than be subject to the time restrictions of regular admission. The room prices are a substantial discount from regular rates: $175 per night Friday, Saturday, or Sunday and $210 per night for Monday through Thursday. Four Seasons Executive Suites are also available for an additional $25 per night.

In addition to the special rates, the spirit of the master can be felt throughout the hotel. Picasso's beloved Mediterranean is the inspiration for a Provencal Sunday Brunch in the acclaimed Fountain Restaurant and dishes with a Spanish taste will also be among the summer's fare. At the entrance to the restaurant, a changing display of real fruits, vegetables and objects will re-create the images that Picasso may have observed as he worked. The exhibition catalog will be open to the photograph of the corresponding painting.

The Four Seasons Hotels is a neighbor of the Museum of Art on Benjamin Franklin Parkway. The brief, picturesque stroll to the museum passes many of the city's other major cultural attractions.

For reservations or information, call, 215-963-1500, or 800-332-3342, or your travel agent.

Flower arranging

The Wynnefield Branch Library, 54th and Overbrook Avenue, will present a two-part series on flowers and herbs. On Wednesday, June 10, Elaine Thomas, Floral Designer at the Flower Gallery will conduct a "Silk Flower Arrangement Demonstration." She will show how you can easily and economically create ambiance both indoors and out with two arrangements — one suitable for a patriotic holiday picnic and the other for an elegant summer dinner.

B Written Comments from the Scholars

Summer Minority Institute
May 20 – June 12, 1992
Final Evaluation

Please answer the following questions as fully as you can on an additional sheet of paper. We value your comments and ideas.

1. Which seminar or module did you find the most informative? Which related closest to your interests? Which talks would you prefer to have longer/shorter?
2. On what topics would you like have a mini-course, like the probability and statistics mini-course?
3. How did you find the C labs? How can these be improved?
4. Would you want more/less *Mathematica* labs?
5. Is the lab work exciting? What other topics might you have chosen for your programming projects? How could the projects be improved?
6. Can you provide some feedback about the lab instructors: Jeff, Libby, Phil, Robbie, and Anthony?
7. Events: Which weekend event did you like the most?
8. Which did you like the least?
9. Mentors: How did you find the mentors' help? Where they there when you needed them? Is there something in addition that they should have done?
10. Course director: Was your interaction with the course director worthwhile? Would you prefer more/less interaction? Was it helpful to eat several meals with the director?
11. Administrative matters: What can we do better next year to improve the institute?
12. Scheduling of courses: Did the sequence of courses make sense? How do you think it could be improved?
13. What was the difference between the interaction with the graduate students and the interaction with the faculty?
14. Have you learned something about how research is done?
15. Were the stipends satisfactory?
16. Overall, was the experience stressful or enjoyable? Please elaborate.
17. Has this experience changed how you think about graduate studies?
18. Additional comments: Please let us know what you think can be done to improve this institute.

Summer Minority Institute

1. The seminar on databases, graphics, longer-graphics & database, shorter- probability & statistics

2. Graphics

3. The 'C' labs were helpful. The approach to teaching 'C' could have been more organized.

4. More Mathematica labs with more things for us to do.

5. The lab work is exciting. I don't know. Make them more related to what we are interested in.

6. Jeff was an excellent teacher. He made things understandable, but he & Libby would go off on tangents & we would get lost to what they are talking about.

Libby was fine, but she always had a new way of doing something without just explaining the old way.

Phil was an excellent teacher too. He went in depth when you didn't understand.

Robbie explained things very clearly & made [illegible] understand what he wanted.

intro was very helpful, but we didn't need too much of his help, because the tutorial was indeed.

7. The weekend horse-back riding & Grad Adventure

8. N/A

9. Their help ~~were~~ satisfactory. Yes, they were. nothing

~~10~~.

10. Yes, it was, more, yes

11.

12. The sequence was fine, but the days were too long.

13. The graduate students gave us the low down on what a 'grad student's' life is really about.

4. Yes, I learned that it's very time consuming.

15. Yes

16. Enjoyable, because I knew that I was not going to be graded or berated for not completing every lab work.

17. Yes, it's made graduate school something more accessible.

18 ① Don't make the lectures for 3 hours long. It becomes boring after awhile.
② Have more fun programs
③ Have books, like for the C class, so students can know what the lecturers are talking about.
④ For the graphics, have more things to do, than just the tutorial on Jack.

Answers to Summer Minority Institute

1) The seminar that I found most informative was Prof. Mertz's seminar on Probability and Statistics. The seminar that related closet to my interest was the one on Artifical Intelligence. The (AI) could have been longer and the Prob. & Stat. could have been shorter.

2). More multi-courses on Graphics & AI.

3) The C labs were very interesting they could be improved by them teaching C and not just a CRASH Course on C.

4) I would want more Mathematica labs.

5) The lab work was exciting and I have no exceptions

6) Jeff - was very nice and helpful.
Libby - Okay but not a very good instructor
Phil - was excellent and a very good instructor who took the time to explain everything.
Robbie - was also excellent and a VERY GOOD instructor &

Willonda Fisher

also took time to explain everything about Mathematica.

Anthony - was very helpful but there was not much he could do because we had a well written tutorial so he just had to point out our mistakes. Of which did a great job.

7) I liked horse back riding the most.

8) Liked them all

9) The Mentors were very helpful & informative.
Ed - Was okay although he acts about 40 but otherwise he was very helpful
Ralph - was GREAT and very helpful.

10) Greg Floyd was very nice and very, very funny. He like to read a lot but that's okay he was helpful.

11) To better improve the institute FIRST we could have met the President of the University, we don't even know his name & no one has mentioned him to us, it is as if we are not important enough to meet him or her.

12) Some of the courses made sense and some because of the scheduling of the courses that are asleep. Make the courses shorter they are too long.

13) The graduate students told us how Graduate School is really is ~~not~~ from their point of view and the prof. told us from their point of view.

14) Yes I have learned a great deal about how research is done.

15) They stipends were satisfactory.

16) The experience has been enjoyable yet long because the days were very long.

17) Yes I have ~~definately~~ definitely decided to go to grad school but have not decided if I will go directly after undergraduate or if I will take some time off

18) Shorten the days & make the weeks longer and also try to give credit for the courses. Also tell the newspaper the Pennsylvanian to stop misquoting the students. I did not appreciate it!

Final Evaluation

1. Mathematica. I found it very interesting to see how probability and statistics are used within this system.

2. I would have preferred to have a mini-course on C programming.
3. I found them to be very informative. I have had a little background in C, so the labs helped to enhance what I knew and I learned more about the language. the labs can be improved by arranging them to be more structured. Some people had much less experience than others, so some people felt confused & DIDN'T KNOW WHAT WAS GOING ON.
4. I would like more Mathematica. I really enjoyed working with this system.
5. Yes.

6. On the whole they were all great. Each one had their own unique way of instructing in the lab. All of them were patient and always willing to help if help was needed. They were very excited about working with us and seemed to take a sincere interest in each student.
7. June 5-7th → movies, horseback riding, Great Adventure, party.
8. I enjoyed them all
9. Ralph and Ed were of great help. They made me feel @ home. They were always open to listening to me and helping me in any way they could.

10. Yes, DR. Provan was a Good director. I Do not feel that there was enough interaction with him. Aside from the seminar he presented, we did not really get the chance to sit down and talk with him in an informal forum as I would have liked. [redacted] Having meals with him was OK., however basically everyone was concentrating on eating, therefore there was not much concentrated conversation.

11. Make sure that all letters verifying acceptance reach the student. Possibly send one to the school address and home address.

12. The schedule was o.K.. I think there should be more balance as far as the time length of the classes.

13. I enjoyed both interactions, The graduate students really put things into perspective. They did not come across as being unrealistic. They were open with us, and answered all of our questions. They were straight-foward and frank. I appreciated them for being upfront and honest— Calming my fears and relieving my curiosities.

14. Yes.

15. Yes.

16. The experience was enjoyable. I could have been @ work now making much money, however money is not my #1 priority; I could not pass up the experience. I am very glad that I got the chance to be apart of the ARO Penn Summer Minority Institute.

17. Yes. I have put things much more into perspective. Before, I was unsure about some things dealing with graduate school, but NOW I Know More and I am better apt to make a GOOD DECISION concerning my continuing education.
18.

Carlos M. Santiago

1. The natural language one and the C related labs and topics.
2. Object-oriented programming; graphics.
3. The C labs were complete and very good. Maybe you can give options, from the start of the program, for the more advanced students.
4. Less Mathematica labs.
5. I liked the C labs and I think they can be more "varied". The project could have several parts that could have been divided between students in a more or less Software Engineering class where you have principal architect, designers, coders, testers, and the kind. Or maybe a project with different options for the more advanced where, for example, one could begin from part II assuming that one knows what was behind part I. I want to see more up-to-date software tools and techniques like object oriented programming, data structures, and abstract data type definitions and implementations.
6. Jeff: I like the way he talked to us with confidence and complete knowledge of what he was saying at all times. He had ideas and was very active in showing us all demos.

   Libby: Very "active". She was real good in her topics and wanted to make us sure that we understood what she was trying to say at all times.

   Phil: Original idea and was very, very serious about us getting the project to run. He was very similar to Jeff and he always had a solution for everything.

   Robbie: He knew very well the Mathematica package and was able to help at all times.

7. The beach.
8. Museums
9. The mentors were both at our side when we needed them. All days they asked us if we were OK and if we needed something. I believed they did a really good job.
10. I think our interaction with him was fine. Yes, it was helpful to eat with him because we had the opportunity to ask him anything in case of doubts or questions.
11. For next year you can divide in a more structure way all topics. I don't believe a course in Statistics was appropiate because I had the idea that this was going to be a more programming oriented internship. I know Statistics are applied but I was expecting a level higher ~~that~~ than that. Maybe you could concentrate half the time in programming techniques of today and the other half time in applications and give the project at that time. Also, be sure of the quantity of money involved in all travel expenses because a trip from PR to U.S. does not costs $350.
12. The sequence of courses was fine.
13. I think the difference is obvious. From one side you get the student position and on the other the professor. But, both were real sincere and very positive.
14. Yes, I have now a complete idea of almost everything that is involved.
15. Yes, they were right.
16. It was enjoyable. Everybody did their best to comply with what they wanted to do.
17. Now I have a complete idea of what I'm going to do.

1. I found computer graphics the most informative.
   Brown's presentation of TramaAid was related closest to my interests.
   I would like a longer talk on networking.
2. C++ objected oriented
3. The C labs were excellent
4. more · Mathematica
5. yes the work is exciting
6. All of the lab instructors were eager to help us and they put in a lot of effort to make sure we understood.
7. African American Art Tour
8. Great Adventure
9. The mentor's were very helpful. At times it seems they would overextend themselves, but it was good to know that they were here to help us it they did it with sincerity.
10. The interaction was worthwhile but I think we need more, not only when we eat but between courses so that there can be a timely exchange of concerns or ideas about the courses.
11. Make sure that everyone involves is informed about all aspects of the program, so that the staff and the students will be expecting the same thing.
12. The sequence made sense.
13. The graduate students and the faculty made us feel at ease, but we could relate to the graduate students better.

14. Yes
15. Yes

16. Yes the experience was enjoyable because the people who were involved were enjoyable people and they wanted us to learn as well as have a good time

17. Not really because I have always wanted to go to graduate school.

1. I found that the C Lab was most beneficial because The instructors were diligent in explaining our tasks and explaining C Language. The discussions ~~were~~ lasted a lengthy time but the information was interesting therefore, I did mind the time I spent in Lecture.

2. –

3. The C Labs were conducive to learning. The labs were always open and clean. The Labwork was challenging and the projects were interesting.

4. more

5. The C labwork was fun and exciting because the projects ~~[redacted]~~ that we worked on were currently being processed in the computer science researching areas.

6. Jeff → Phil, Robbie, and Anthony were all very, very enthusiastic ~~[redacted]~~ in helping us learn, create and implement our particular programs, and excercises. It felt very comfortable asking them questions. Libby was excited about explaining our projects but was not around during our projects

7. I liked the horseback riding the most.

8. N/a

9. My mentors (Ralph and Ed) were very easy going ~~[redacted]~~ but they took care of business on time.

10. I would like to have had at least one day free to ~~[redacted]~~ visit a certain professor.

11. N/a

12. yes

13. There was more interaction with the graduate students and less interaction with the faculty. I appreciated the time that we were alloted to converse with the graduate students. But, I would have liked and enjoyed more time to talk

with the professors.

14. very little

15. Yes

16. enjoyable. I learned about graduate school, studies I probably would not have seen until I was a Junior or Senior in College, and I have been exposed to many different aspects of Computer Science.

17. Yes - In a very positive way.

18. I loved this program.

1) THE WORK ON UNIX WAS THE MOST INFORMATIVE & INTERESTING. I FELT THE LECTURES AFTER LUNCH COULD BE CUT DOWN.

2) GRAPHICS GENERATION

3) THE C LABS WERE GOOD, BUT A CONSULTANT COULD BE HELPFUL.

4) MORE

5) THE PROJECTS WERE WELL STRUCTURED & ENGAGING. NO CHANGE RECCOMENDED.

6) Jeff - informative & knowledgable, get ahead too far sometimes
   Libby - was around enough to comment
   Phil - congenial and and excellent consultant & lecturer
   Robbie - willing to help and adaptable to student's queries
   Anthony - competent & reliable

7) Horseback Riding weekend

8) Some city trips (too much walking)

9) There were times when Ed seemed ~~asse~~ inaccessible, but those were very few & in general, he and Ralph were extremely helpful.

10) Dr. Provan was very understanding and patient and accomadating, while also making a definite effort to keep us informed and to provide some guidance along the course. Eating w/ him provided opportunities to discuss future plans & current complaints

11) Better organization (schedules, etc) & more money would help.

12) Course order was well planned and flowed together.

13) The grad students were more informal and willing to give opinions on professors & their respective abilities

14) Yes.

15) Yes.

16) Enjoyable overall, but there was some stress related to the extensive 2 63 hr. lectures (I would get tired), and the environment of the city is not the most settling. But one adapts.

17) No.

18) Make it longer (maybe 5 weeks) if possible. Other suggestions are above.

1. Most informative: C Programming, Probability and Statistics Mini-course
Closest to my interests: Databases, C Programming
Increase: Databases, Linear Algebra
Decrease: Probability and Statistics (or split up sessions if possible)

2. Databases, if more ~~in~~ ~~it~~ detailed than the databases we worked with.
The probability and statistics mini-course should be repeated.

3. The C Labs were ~~in~~ fairly adequate. They can be improved by increasing the area of lab, perhaps adding a blackboard, and making sure that all of the students ~~bot~~ who want/need terminals have them.

4. I would like the same number of Matematica Labs but increase the difficulty of the assignment from just typing in code and doing a function that Dr. Mirsky showed us, to perhaps trying to obtain specific results.

5. The final project was interesting. The projects can be improved by being distributed earlier in the program and giving the students choice to choose from, ~~and make the projects~~.

6. Jeff - provided insight on the "real" world, helpful
Libby - know her stuff but not always sure, often 'bickered' with Jeff on how to do things
Phil - ~~the~~ the best teacher of the lab instructors, very energetic
Robbie - nice guy, helped explain what the assignment was when you did not understand.
very knowledgable of Matematica

7. Horseback riding and Great Adventures and tours of Philadelphia

8. Horseback riding

1. The mentors helped when they could and helped us locate stores and places of interest in the area. For the most part but at times they were hard to get in contact with but they were around during the crucial moments. They could have been more coordinated between themselves and at more informed and sure of what was 'supposed' to happen.

20. Yes. More. Yes

1. If possible increase the number of students and the length of the program.
2. Make sure plans are furnished

2. The scheduling could be improved by not scheduling the one lecture professors at the end of the program.

3. The graduate students were more approachable ~~but~~ The faculty was helpful and answered every question if they could, in a friendly manner.

1. Yes.

2. A little more would help.

3. The experience was enjoyable. I met several nice people with whom I hope to continue correspondence. I learned more about graduate school and research. I also saw Philadelphia.

1. Yes.
2. Give the students a choice of what they want to study or have the program and the program description be more specific and orientated towards an objective.

G. TSATSOY.

① Mathematica was the most interesting part of the seminars, because I can see how it could speed up & make more efficient the research done at my school.

② MORE IN DEPTH MATHEMATICA APPLICATIONS.

③ THE C-LABS COULD BE MORE EFFICIENT IF ONLY ONE INSTRUCTOR SPOKE ON A TOPIC AT A TIME.

④ MORE.

⑤ THE PROJECTS COULD BE IMPROVED IF THERE WERE TWO PROJECT GIVEN TO THE STUDENTS. ONE BASED ON ABILITY, AND ONE AS A GROUP PROJECT.

⑤ THEY ARE ALL VERY KNOWLEDGEABLE AND EFFECTIVE IN THE LABS.

⑦ GREAT ADVENTURE.

⑧ THE PARTY IT COULD HAVE BEEN DONE BETTER.

⑨ THE MENTOR WERE EXCELLENT.

⑩ EATING @ THE DIRECTOR HELPED BECAUSE IT INSURED THAT WE WERE GETTING ENOUGH INFORMATION W/O IT BEING TOO MUCH.

⑪ MAKE IT LONGER. LESS CSI ORIENTED.

⑫ YES; YES MAKE IT MORE VARIED. IF YOU ARE TRYING ATTRACT ALL DISCIPLINES TO GRADUATE SCHOOL, THEN THERE SHOULD BE SOMETHING HERE FOR EVERYONE.

(13.) No much.

(14) Yes.

(16). Enjoyable. It allow us to absorb what we could w/o having to push us over the edge.

(15) Yes.

(17) Yes.

(18) The institute needs to be a little more organized.

1. The seminars that I liked the most were graphics and Databases. You should give more time to Dbases an less to probability.
2. I think that a mini-course in graphics with lab sections would be more interesting.
3. C Labs were good but you can improve then by separating the more experienced people from then less experienced ones and give each group programs according to their level and skills.
4. If you reduce the Prob-Stat course to the really important things then the Mathematica Lab would be enough to do whatever you want to do.
5. The best part of the SMI were the lab sections. Improvement can be done by giving each set of persons a theme in the programing level that they are. Interesting topics could be nathural language and graphics.
6. They were all very nice and helpfull. Wonderfull sellection.
7. I enjoyed the most the last weekend. Great Adventures and horseriding.
8. All were good.
9. Another wonderfull selection.
10. Greg Provan did an exelent job since the very beging. The interaction with him was good and also the meals.
11. Money was god but I think that you shoud require assistance to all lecturs to make sure that the money you are paying is well used.

12. Schedual e was almost good but please don't start at nine. Besides that everything was good.
13. students tell you the real thing about grad shools. Faculty members are good but they hide many details that students should know.
14. I thought that research was more personal but I found something interesting about teamwork.
15. I think the stipends were enough.
16. I really like the experience but I was expecting something different. I do not mean that I did not liked but just that it was different.
17. I thought that grad school was shoter (in terms of time). You have to really think hard about going to it.
18. The preceding answers tells what I think can be done to improve the program. I just want to add thank you for programs like these and keep doing things like this.

1. The seminar I enjoyed the most was Parallel Processing with Professor Rajas Keratam. I enjoyed the lecture on the Connection machine and the discussion of the concepts of parallel processing with ALU modulation.
I would have preferred more time with Database Management

2. Natural Language
C++

3. They were very up to date and very fast.
They can be improved by having more tutors in the evening hours

4.— More Mathematica labs.
5. Very exciting. The programs covered most topics well
6. Jeff- Very down to earth he explained the topics well.
Libby- She was motivating and very energetic
Phil- I liked the way he explained pieces
Robbie - He knew Mathematica very well.
7. The weekend where we went to Horseback riding.
8. None
9. Mentors,
ED & Ralph were very alert and responsible
They were there when we needed them.
They could have made a list w/ everybody's phone # in the program.
10. More interaction with the director, he was very helpful
11. See if the students can get sent a ticket (round trip) before coming to the program.
12. The courses format was very well prepared
13. I enjoyed the evening lyceums & dinner discussions
14. I learned that research is not the most glamorous life but it is rewarding and challenging.

15. The stipends were excellent.

e. The first days were a little cold, but once the weather warmed up everything was fine. I enjoyed the constant challenge of the courses and all the cooperation that we received from the instructors & T.A.'s. I enjoyed the opportunity to get a "hands-on" experience of the computer systems.

1. I have always intended on going to graduate school.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	July 1995	Final 1 Oct 89 - 14 Mar 95

4. TITLE AND SUBTITLE
Artificial Intelligence (AI) Center of Excellence at the University of Pennsylvania - Final Report

5. FUNDING NUMBERS
DAAL03-89-C-0031

6. AUTHOR(S)
Aravind Joshi (PI), Ruzena Bajcsy, Norman Badler, Peter Buneman, Bonnie Webber (co-PIs)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Department of Computer and Information Science
University of Pennsylvania
200 S. 33rd Street
Philadelphia, PA 19104-6389

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The ARO AI Research and Education Grant was initially awarded to the University of Pennsylvania in June 1984, and renewed in October, 1989. The AI Center was established to integrate and coordinate our various research efforts in AI in which we have excellent strengths. The key thrusts were in the following areas: 1) Natural language processing: language and speech; 2) Machine perception and robotics; exploration and perceptual development; 3) Task oriented computer animation; 4) Programming structures for databases and knowledge bases; 5) Parallel processing in Artificial Intelligence. This final report outlines our research in all the above areas, more specifically in Computer Graphics and Animation; Natural Language Processing; Diagnostic Reasoning and Expert Systems; Computer Vision and Robotics; Neurally Motivated Models; Programming Language Design and Theory; Heterogeneous Databases; Real Time Systems; and Algorithms in Computational Biology. This final report also includes the final reports for each of the sub-contracts awarded under the AI Center. The final reports for the sub-contracts are found in Appendix B.

14. SUBJECT TERMS
Artificial Intelligence, natural language processing, machine perception and robotics, task oriented computer animation, programming structures for databases & knowledge bases, parallel processing in AI

15. NUMBER OF PAGES
154

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL